UNIVERSITÉ PARIS 1
PANTHÉON SORBONNE

THÈSE pour l'obtention du titre de Docteur en Sciences Économiques

Sur les effets économiques locaux des événements extrêmes dans les pays à revenu faible et intermédiaire

Présentée par

Robert Max Reinhardt

Directrice de These: Lisa CHAUVET
Professeur des Universités, CES, Université Paris 1 Panthéon-Sorbonne

Membres du jury: Katrin MILLOCK (Presidente)
Directrice de Recherche CNRS, Paris School of Economics

Raja CHAKIR (Rapporteur)
Directrice de Recherche INRAE, AgroParisTech, Paris-Saclay

Stéphane GOUTTE (Rapporteur)
Professeur des Universités, Université Paris-Saclay, UMI SOURCE- IRD - UVSQ

Paolo AVNER (Examinateur)
Économiste principal, Facilité mondiale pour la prévention des catastrophes et le relèvement (GFDRR), Banque mondiale

Kenneth HOUNGBEDJI (Examinateur)
Chargé de Recherche IRD, Université Paris-Dauphine & PSL, LEDa-DIAL

Andrew J. KRUCZKIEWICZ (Examinateur)
Chercheur, National Center for Disaster Preparedness (NCDP), Columbia Climate School, Columbia University

Centre d'Économie de la Sorbonne (CES)
École Doctorale D'Économie (ED465)
Université Paris 1 Panthéon-Sorbonne

Soutenu publiquement le 23 Septembre 2025

UNIVERSITÉ PARIS 1
PANTHÉON SORBONNE

THESIS for the title of Doctor of Philosophy in Economics

On the Local Economic Effects of Extreme Events in Low- and Middle Income Countries

Presented By

Robert Max Reinhardt

Supervisor: Lisa CHAUVET
University Professor, CES, Université Paris 1 Panthéon-Sorbonne

Thesis Committee: Katrin MILLOCK (President)
Research Director at CNRS, Paris School of Economics

Raja CHAKIR (Reviewer)
Research Director at INRAE, AgroParisTech, Paris-Saclay

Stéphane GOUTTE (Reviewer)
University Professor, Université Paris-Saclay, UMI SOURCE- IRD - UVSQ

Paolo AVNER (Examiner)
Senior Economist, Global Facility for Disaster Reduction and Recovery (GFDRR), World Bank

Kenneth HOUNGBEDJI (Examiner)
Research Manager at IRD, Université Paris-Dauphine & PSL, LEDa-DIAL

Andrew J. KRUCZKIEWICZ (Examiner)
Researcher, National Center for Disaster Preparedness (NCDP), Columbia Climate School, Columbia University

Centre d'Économie de la Sorbonne (CES)
École Doctorale D'Économie (ED465)
Université Paris 1 Panthéon-Sorbonne

Publicly defended on 23 September 2025

This page is intentionally left blank.

To my family.

Acknowledgments

Writing and completing this thesis has been a challenging, rewarding, and intellectually enriching journey. When I began my Master's at Paris 1 six years ago, I could not have anticipated the many academic and personal adventures that were to come. Fortunatetly, I was never alone in them, thus I owe my deepest gratitude to the wonderful colleagues, friends, and family who supported, encouraged, and helped me throughout this adventure. The number of people who have supported me along the way is too great to name individually, and I can only hope no one feels overlooked.

First and foremost, I would express my sincere gratitude to my supervisor, Lisa Chauvet, who could not have done a better job to guide me. Your guidance, feedback and advice have always made me feel sure about the paths I was following. Despite my partially chaotic stream of thoughts of new climate models, you always patiently listened, gave me the boundaries I needed to advance and helped me to not lose my sight. You taught me the importance of maintaining focus and striving for rigor in my research with your kind and supportive guidance. For that, I am deeply grateful—thank you!

I am also honored to have Katrin Millock, Raja Chakir and Stephane Goutte as leading members of my jury, as well as having Kenneth Houngbedji, Paolo Avner and Andrew Kruczkiewicz as external members. All your comments were and will be extremely helpful in improving the quality of this thesis and in shaping me as a researcher with an interdisciplinary focus. Thank you very much.

Next, during my thesis I had the privilege to spend time at the University of the UN in Bonn, Germany. I would like to thank Maxime Souvignet and his great colleagues for the warm welcome and the fruitful discussions on climate adaptation modelling and impact modelling. Equally, I would like to greatly thank Scott Barrett at Columbia University for supporting my visit in New York. Despite the short stay, I had the privilege to exchange with wonderful, smart and future-oriented researchers that motivated me to pursue my interest in disaster modelling. I would like to thank the professors, PhD students, and postdoctoral researchers at SIPA and the Climate School who warmly welcomed me and helped me focus on pressing climate-related issues. I am also grateful to the staff at NASA-GISS with whom I had the opportunity to engage. Also my thanks here to the Région Île-de-France, the Alliance Program and the Centre d'Économie de la Sorbonne (CES) for financial support during these visits. A final thanks to Giovanni, Lucile, Julius, Gaspard and all others that went with me to

New York and made the time special.

I would also like to extend my gratitude to the professors at Paris 1 Panthéon-Sorbonne at the CES, who provided valuable feedback in internal seminars or informal exchanges. Especially, I would like to thank Remi Bazillier, Clement Bosquet, Morgane Laouenan, Laurine Martinoty and Clement Gorin. A special thanks also to Morgan Hull Brousmiche, who helped me to get settled and supported me throughout the three years of teaching at the PSME program.

This journey would also not have been the same without the various generations of PhD colleagues and friends at the Maison des Science Economiques (MSE). At the start of the PhD, everything felt overwhelming. I am deeply grateful to Yasmine, Thibault, Andrea, Jala, Fabio, Liza and Stephen for being there from the beginning — offering encouragement, friendship, and a warm welcome into the PhD journey. Thanks to Adham, always being a wonderful companion for coffee and chatting about disasters and being a pleasant co-author to work with. From the very beginning, I was privileged to be also surrounded by amazing colleagues and wonderful people— becoming now close friends: Maria, Farida, Nada, Solèn, Bea, Rosa, Rengim and of course my friends from the Master Johann and Carlos, thank you for all the laughs and good moments that always kept me smiling throughout the four years. I will always cherish the great moments in and outside MSE. My experience would not have been the same without you. To Antonio, who joined me in the Banquier office and kept me wonderful company in gray research days with his smiley attitude, thank you. Also my very big thanks to the new generation of inspiring economists - Carolina, Abigail, Alejandro, Yanis, Afifi - it was great having you around, taking lovely coffee breaks, playing Tradle and from time to time being motivated to also go out of the office to do some joint fun activity. Last but not least, I'm truly grateful to Maxence, Coralie, Olivier, Hannes, Hugo, and Pol—not just for always being around, but for the great time and the camaraderie. It was a real joy to share this journey with you.

Outside of work, I was fortunate to have kind and loving support from people around me that I hold very dear, even if they were not in Paris. To Magdalena, Jenny and Linden thank you for the emotional support and for always being there. I am deeply grateful to have you in my life. I would like to extend my profound gratitude to my parents, I would not be here today without you, without your continuous support and encouragement throughout the years. Finally, I would like to extend my thanks to my partner, Caroline, who throughout the masters and the PhD was always on my side with encouraging words, patience and love so that I never felt alone in the process.

This page is intentionally left blank.

On the Local Economic Effects of Extreme Events in Low- and Middle Income Countries

Abstract

Climate change is one of the biggest challenges of the 21^{st} century. As many regions around the world experience greater shifts and fluctuations in weather patterns, which in some cases lead to more extreme events, it becomes more important than ever to study their impacts. This is particularly true for low- and middle-income countries, where less economic resilience can be found. This thesis aims to contribute by highlighting the complexity of hazards that impact economic systems in low- and middle-income countries. Extreme events, especially destructive ones, not only destroy houses and cause fatalities but can also propagate through the economy via supply chains. This is discussed in Chapter 1 using an event-study difference-in-difference model for the case of foreign investment inflows and earthquake exposure in Indonesian districts. The chapter shows not only the direct effects in affected districts. It also documents indirect effects that occur when connected industries, serving as suppliers, are impacted. Many regions are exposed to more than just a single type of hazard. Thus, to better understand hazard environments, Chapter 2 introduces a new compiled data set. In this chapter a newly constructed global gridded dataset is introduced, creating a monthly global 10x10km dataset for 17 hazards and five socioeconomic outcomes. Finally, based on the improved data access, the thesis revisits well-studied economic literature strands from a multihazard perspective. First, Chapter 3 examines how extreme events, particularly those that co-occur, impact population changes in rural Sub-Saharan Africa by employing a panel three-way fixed effects model. Categorizing hazards as binary dummies can obfuscate their actual effects. Instead, cumulative shocks within a year and the overlapping of events cause fundamentally different responses to shocks that should be considered as "new" events. Concerns for research of spatial spillovers, the co-occurrence of events, dynamic effect carry-overs and non-linearities are discussed. Chapter 4 shifts the focus to urbanization in Sub-Saharan Africa, utilizing a panel model with three sets of fixed effects. Weather shocks are found to harm urbanization in the short term. This is particularly the case for floods, whose effects are amplified if floods follow heatwaves. The article also highlights that the spatial scope of a shock matters and historical treatment matters, especially for droughts. All of these articles emphasize the importance of focusing in economic studies on the different causes and impact dimensions that extreme events have on economies. The findings also have high policy relevance, as they emphasize in all chapters that rural, more vulnerable economic agents are especially affected by the effects of extreme events.

Keywords: Climate Change, Development, Economic Growth, Environmental Economics

Résumé

Les effets du changement climatique constituent l'un des plus grands défis de notre époque. L'une des principales manifestations de ce phénomène est l'augmentation du nombre de chocs climatiques ces dernières années. A ce titre, cette thèse de doctorat cherche à quantifier les conséquences sociaux-économiques de différentes intempéries en tenant compte de leurs complexités, mais également des interactions existantes entre divers évènements météorologiques. Les effets sur les pays en développement sont particulièrement importants, car ils sont les plus vulnérables en raison de leur capacité de réaction relativement limitée et de leur exposition fréquente à des phénomènes climatiques extrêmes. Le chapitre 1 examine les canaux de transmission d'un type spécifique de chocs, en prenant l'Indonésie et les flux d'investissements étrangers comme étude de cas. En utilisant la méthode de la double différence, l'article montre que les tremblements de terre n'affectent pas seulement les districts, mais qu'ils ont des effets sur l'ensemble de l'économie en raison des liens industriels. De nombreuses régions connaissent plus d'un type d'aléa. Pour mieux comprendre la répartition d'aléas, le chapitre 2 présente un nouvel ensemble de données. Il s'agit d'une base de données mondiale mensuelle de 10x10 km pour 17 aléas et cinq variables socio-économiques, construite à partir d'ensembles existants. Sur cette base de données améliorée, la thèse réexamine des pans bien étudiés de la littérature économique dans une perspective multirisque. Le chapitre 3 examine l'impact des événements extrêmes — en particulier lorsqu'ils se produisent simultanément — sur l'évolution démographique dans les zones rurales d'Afrique subsaharienne. En utilisant un modèle de panel avec trois effets fixes, l'article montre que la catégorisation des risques en tant que variables binaires peut masquer leurs effets réels. À la place, les chocs cumulatifs sur une année et le chevauchement d'événements entraînent des réponses fondamentalement différentes, à considérer comme des événements « nouveaux ». Des préoccupations liées aux retombées spatiales, à la cooccurrence, aux effets dynamiques reportés et aux non-linéarités sont également abordées. Le Chapitre 4 porte sur l'urbanisation en Afrique subsaharienne, en utilisant un modèle de panel avec trois séries d'effets fixes. Les chocs météorologiques nuisent à l'urbanisation à court terme, en particulier les inondations, dont les effets sont amplifiés lorsqu'elles suivent des vagues de chaleur. L'article souligne aussi que la portée spatiale d'un choc est importante, surtout pour les sécheresses. Tous ces articles soulignent l'importance de se concentrer, dans les études économiques, sur les différentes causes et dimensions de l'impact des événements extrêmes sur les économies. Les résultats ont une forte pertinence politique, car ils montrent dans tous les chapitres que les agents économiques ruraux, plus vulnérables, sont particulièrement touchés.

Mots Clés: Changement climatique, developpement, croissance economique

Contents

List of Figures

List of Tables

List of Acronyms

ASEAN	Asian Association of Southeast Asian Nations
CAPE	Convective Available Potential Energy
CID	Climate Impact Driver
CP	Convective Precipitation
DFO	Dartmouth Flood Observatory
EM-DAT	Emergency Events Database
ENSO	El Niño-Southern Oscillation
FDI	Foreign Direct Investment
FEMA	Federal Emergency Management Agency
GDP	Gross Domestic Product
GHSL	Global Human Settlement Layer
GISA	Global Impervious Surface Area
HDI	Human Development Index
IPCC	Intergovernmental Panel on Climate Change
MJO	Madden-Julian Oscillation
MMI	Modified Mercalli Intensity
MRIO	Multi-Regional Input-Output Table
NDVI	Normalized Difference Vegetation Index
NOAA	National Oceanic and Atmospheric Administration
NTL	Nighttime Light
SDEM	Spatial Durbin Error Model
SDG	Sustainable Development Goals
SPEI	Standardized Precipitation Evapotranspiration Index
SREX	Special Report on Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation
SSA	Sub-Saharan Africa
SUTVA	Stable Unit Treatment Value Assumption
TWFE	Two-Way Fixed Effects
UN	United Nations
USGS	United States Geological Survey
VEI	Volcano Explosivity Index
WBGT	Wet-Bulb Globe Temperature
WMO	World Meteorological Organization

This page is intentionally left blank.

General introduction

Anthropogenic climate change influences not only the climatic world in which we live, but also especially the interactions humans have with the nature surrounding them. Climate change, classified as one of the nine planetary boundaries by Rockström et al. (2009), represents a critical threshold within which humanity can operate safely. As one of the most significant and widely discussed challenges of the 21st century, it underscores the need for global cooperation to maintain the stability of the Earth's system. Failure to do so not only poses a threat to economic prosperity, but also puts the achieved development goals, for example, captured with United Nations (UN) Sustainable Development Goals (SDG), at risk. Since the discovery of greenhouse gases by the works of Fourier, Tyndall and Arrenhius in the 19th century, an increasing number of studies and articles have been brought forward in the mid-20th century since the establishment of modern climate science (for an overview, see Dobes, Jotzo, and Stern, 2014). Since then, the world community has established various international agreements, notably the Kyoto Protocol and later the Paris Agreement, and joint working projects such as Intergovernmental Panel on Climate Change (IPCC) or the Green Climate Fund to limit the extent and inform about adverse climatic effects. In media and scientific literature, the discourse on climate change is often framed dichotomously: focusing either on mitigation, which aims to limit the extent and speed of global warming, or adaptation, which addresses adjustments to its consequences. This thesis will focus on the latter, especially in a developing context.

Climate adaptation is an umbrella term that encompasses various strategies, including the development of resilient infrastructure, agricultural systems and healthcare systems. Economists have played a crucial role in the scientific debate on such strategies, particularly in the cost-benefit analysis of policies and the assessment of climate-related economic risks. It is of utmost importance to understand the risk that a person, firm, region or country faces from climate change when designing and evaluating adaptation strategies. Following the IPCC on climate extremes, such climatic risk is usually composed of three components: *exposure* ("is an individual/area exposed to an event?"), *hazard* ("how likely is it that an event occurs?") and *vulnerability* ("If so, how vulnerable is a region or population where the event occurs?"). (cp. IPCC, 2012, p.69). Sometimes, the notion of *response* ("How do individuals/areas respond to the

event?") is also added. This thesis will, therefore, center around particular elements of climate risk. Especially the articles combine the notion of exposure by spotlighting the frequency of extreme events and vulnerability by focusing on some of the world's poorest regions. According to the IPCC, there is high confidence within the scientific community that extreme climatic events (such as tropical storm strength or rainfall variability) are likely to increase globally, albeit to varying extents in different geographic areas. Lower- and middle-income countries are disproportionately affected due to higher financial constraints and lower resilience, resulting from higher levels of poverty, poor governance and conflict. An illustration of the statement is shown in Figure 1. Figure 1 shows the global distribution of exposure to extreme events by the human development of each 10x10km cell in the period 2004 to 2015[1]. The data is based on Chapter 2 of this dissertation combined with data on the Human Development Index (HDI). On average, red areas represent regions with a lower Human Development Index compared to the rest of the world. It becomes clear that many parts of the Global South, especially Africa, South Asia, and certain regions of South America, appear to have relatively low levels of HDI but higher levels of observed exposure to extreme events. This might reflect some form of lack of adaptive capacities of households, but also potentially institutions, that varies significantly around the world. Given that the poorest households are already more vulnerable to climate change (see, e.g. Hallegatte, Vogt-Schilb, Rozenberg, et al., 2020), it becomes clear why this thesis places particular emphasis on low- and middle-income countries.

However, the question arises of how to connect extreme events with economic outcomes. To connect climate conditions with outcomes associated with economics and other disciplines, the IPCC introduces the concept of so-called Climate Impact Driver (CID) (Seneviratne, Nicholls, et al., 2012) - specific climate components that directly affect both human systems and ecosystems. These CIDs are usually assessed at a local or regional level, where the interactions between humans or ecosystems and the climate conditions are most pronounced. The concept of climatic conditions as well as extreme events affecting socio-economic outcomes is found in all chapters of this thesis. Although often not referred to as climate impact drivers, many articles at the forefront of economic research have used them, especially as a treatment variable. Notable examples could be the usage of temperature and precipitation ("rainfall") on out-

[1] For the definition of extreme events the following thresholds are taken: heatwave (WBGT>28), flood (DFO Index>1), wildfire (>20% burned), tropical storm (>33 m/s), wind storm (>30 m/s), hail (>1 day with Prob>90%),earthquake (>5 MMI), dust storm (>1 day),landslide (count>1), tsunami (wave >0m), volcano (VEI>0), air pollution (>75 $\mu g/m^3$)

Fig. 1. Global Distribution of Exposure to Extreme Events by Human Development Index (2004-2015)

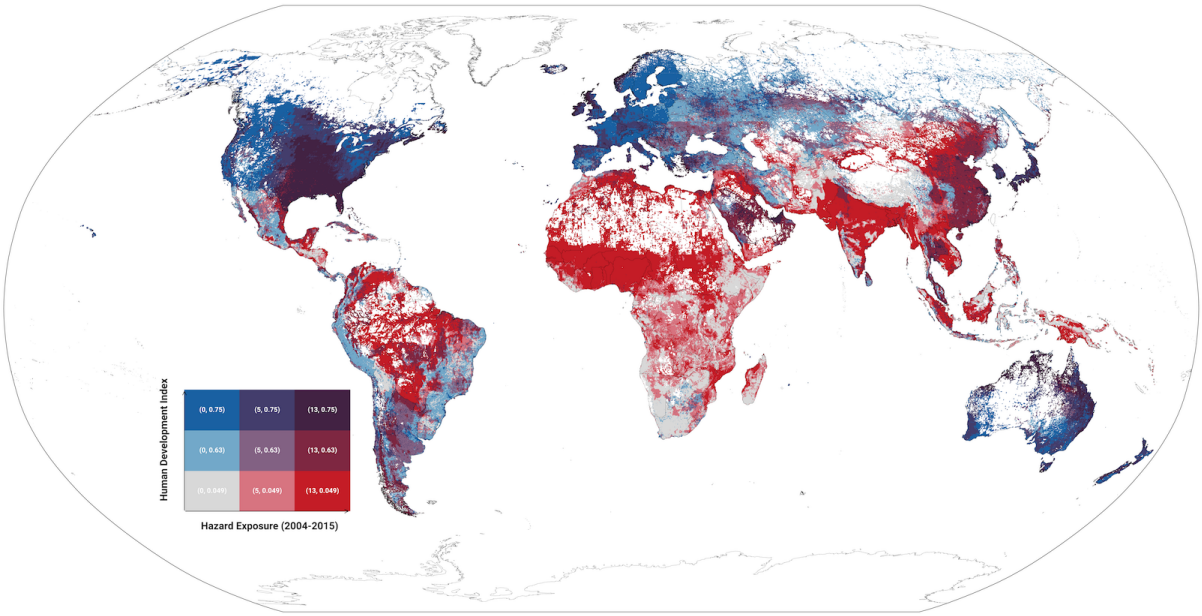

Note: This figure is based on the compiled dataset in Reinhardt (2024a) based on the accumulation of 19 common hazard-types and data on Human Development Index by Sherman et al. (2023). The figure shows the global distribution of exposure to extreme events in a bi-variate plot against the relative Human Development Index (HDI). In white, the specified cut-offs are shown. Only areas with more than 10 people per square km are included.

come dimensions such as crop systems (Aragón, Oteiza, and Rud, 2021; Deschênes and Greenstone, 2007; Lobell, Schlenker, and Costa-Roberts, 2011), cities and settlements, and the connected topic of migration (Beine and Jeusette, 2021; Bohra-Mishra, Oppenheimer, and Hsiang, 2014; Cattaneo, Beine, et al., 2019), health (Meierrieks, 2021; Watts et al., 2017) or more generally on economic growth and production (Nordhaus, 2019; Tol, 2018; Dell, Jones, and Olken, 2012; Burke, Hsiang, and Miguel, 2015). Many of these studies have been cited hundreds or thousands of times. Other studies have focused rather on the effects of the individual type of climatic shocks on the microeconomic dimension of households in firms. Recent notable examples focus, for example, on the impact of wildfires on mortality (Burke, Driscoll, et al., 2021) or the effect of tropical cyclones on firm assets (Pelli et al., 2023).

A common theme in these studies is that they generally focus on a specific type of climatic, weather or hazard shock, such as temperature, tropical cyclones, or rainfall anomalies, to measure the impact on a specific outcome dimension. Only a few em-

pirical studies explicitly consider the interactions of different types of climate impact drivers or, generally speaking, extreme events (see, for example, Fezzi and Bateman, 2015). In theoretical models, such as the Integrated Assessment Models, interactions are incorporated. However, these often still fall short of accounting for the full complexity of interactions. Simplifying the impact of extreme events on outcomes, especially in applied economic articles, to isolated events allows researchers to reduce the complexity of hazard impacts and better establish causal connections. However, insights from climate scientists, earth scientists, or meteorologists reveal that such simplification often ignores the complex nature of the climate system. The literature on extreme events in climatic science has shifted toward adopting a more holistic view of the complex interactions between individual extreme phenomena. In particular, the notion of "connected" or "compounding hazards"[2] has recently seen a surge in interest. This new direction of literature is important for the economics sphere. Focusing solely on individual hazards without accounting for overlapping shocks can lead to an underestimation of general risk or attributing impacts to the wrong event. Thus, this new direction is also crucial for economic policy design. In the following paragraphs, the term and its importance for this thesis will be explained. Throughout this dissertation, each chapter presents the concept that individual risks do not solely influence socioeconomic outcomes. Individual hazards tend to interact with a range of other extreme events or pre-existing vulnerabilities. This thesis encompasses a broader range of hazards than is typically discussed, yet the list of hazards considered is not exhaustive.

Connected Events, Their Importance for the Economic Literature and Methodological Challenges

Before discussing the methodological challenges associated with connected hazards, a brief, stylized introduction is provided on what connected or compound events are and why they should concern economists.

According to the IPCC Special Report on Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation (SREX), compound events refer to

[2] In this thesis, compounding and connected shock or events will be used synonymously. It needs to be acknowledged, however, that "compound" is technically a narrower definition in the climate science literature.

climate-related or weather events that manifest simultaneously, shortly after each other or in spatial proximity to one another and tend to lead to amplified effects of individual events (cp. Seneviratne, Nicholls, et al., 2012, p.118). In addition, using the typology by Zscheischler, Martius, et al. (2020), compound events can be categorized into pre-conditioned events (one climate-driven condition amplifies the impact of another), multi-variate events (two or more climate-impact drivers or hazards happen simultaneously), temporally compounding events (two or more events occur in a short sequence after another) or spatially compounding (two or more events occur in proximity to one another). Each different interaction pair between the type of event and the compound type might have fundamentally different economic impacts. An example is the preconditioning of drought and heatwaves, which increases the occurrence of wildfires (as mentioned in AghaKouchak et al., 2020). Another example is the combination of tropical storms that make landfall, causing not only direct damage through the wind but also damage through flooding and storm surges, which particularly affects coastal communities. In this dissertation, particularly in Chapter 3 and 4, "compound" events are defined as two different events overlapping within a narrow one or two-month window.

There are two notable reasons why it is important to consider connected effects in economic studies: first, the response to such shocks might be fundamentally different from those of isolated hazards. For example, studying the response of temperature anomalies on agricultural production, as is commonly done, could be a good approach to measuring the response function of temperature on soil production or crop outcome. However, if the temperature anomaly occurs simultaneously with a heat wave event[3], the predicted loss in crop outcome found can only be partially attributed to the anomaly in itself but also to the reduced labor productivity of the harvesting farmer. Another example would be estimating recovery periods after a shock. If one presumes only a singular shock took place, the derived policy implications for post-shock recovery might be misguided. Building on the previous example, one might suggest a more effective irrigation system in response to a temperature anomaly, but one would also need to consider strategies to mitigate the heat exposure of farm workers. Second, if an economic study is conducted at the regional or even district level, the variation in exposure to various hazards may differ within the same geographic area. One village might see a flood and a heatwave, while another village 30km away, but in the same

[3] Both are strongly connected, but a heatwave event adds additional environmental factors such as air moisture and wind speed.

district, might experience a landslide instead of a flood due to excessive rainfall. If these other events are excluded, it may cause important variables to be overlooked in the estimation, resulting in omitted variable bias. This is especially a concern for cross-sectional or panel models that include climate variables (Hsiang, 2016). It is, therefore, essential to adopt a multi-hazard perspective, particularly when conducting policy-relevant research.

To emphasize this more, as an example, Figure 2 plots the number of distinct hazard types (e.g. flood and heatwave) recorded between 2004 and 2015 on a 10x10km scale.[4] The data is based on the data from Chapter 2. The figure plots only eleven hazards, including drought, heatwaves, floods, earthquakes, tropical cyclones, wind extremes, landslides, tsunamis, volcanoes and wildfires. The list is non-exhaustive and could be extended across more dimensions, but for illustrative purposes, it shall suffice here. The two bar graphs show the fraction of GDP (left) and population (right) that is exposed to the number of hazard classes[5]. On a global level, the bar plots reveal that for the eleven hazards shown, less than 7% of GDP and less than 6% of the population have not seen any extreme event over the period of 2004 to 2015. Around a quarter has seen a single event. In these areas, it is possible to conduct event studies easily with a single-hazard type treatment. For all remaining 70% of the sample, more than one hazard class occurred, which might contaminate any possible regression model, if not accounted for. This concerns especially the Southern-Eastern parts of the United States, the Sahel Zone in Africa and parts of South-East Asia, including Vietnam, Cambodia and Laos. Thus, the more different hazard classes from this non-exhaustive list occur in a given location, the more the effects in studies using only a single event might be driven by hidden hazards or interactions between different hazards.

When it is so essential, the question arises as to why more research has not included the notion of compounding shocks. There are both practical and empirical reasons for this. It would be a fallacy to assume no research is conducted on the topic. Work in interdisciplinary fields such as the ones by Ridder et al. (2020) and Thalheimer, Choquette-Levy, and Garip (2022) as well as works accounting for accumulative shocks (e.g. Di Falco et al., 2024) over time do study the question of connected shocks. Yet, practically speaking, there was a lack of availability of homogenized datasets that allow for the

[4] The figure does not show the quantity of how *often* hazards occur, just how many different types occurred between 2004 and 2015.

[5] The data is based on computed shares of GDP from Kummu, Kosonen, and Masoumzadeh Sayyar (2025) in 2015

Fig. 2. Illustration of Hazard Presence and Relative Importance

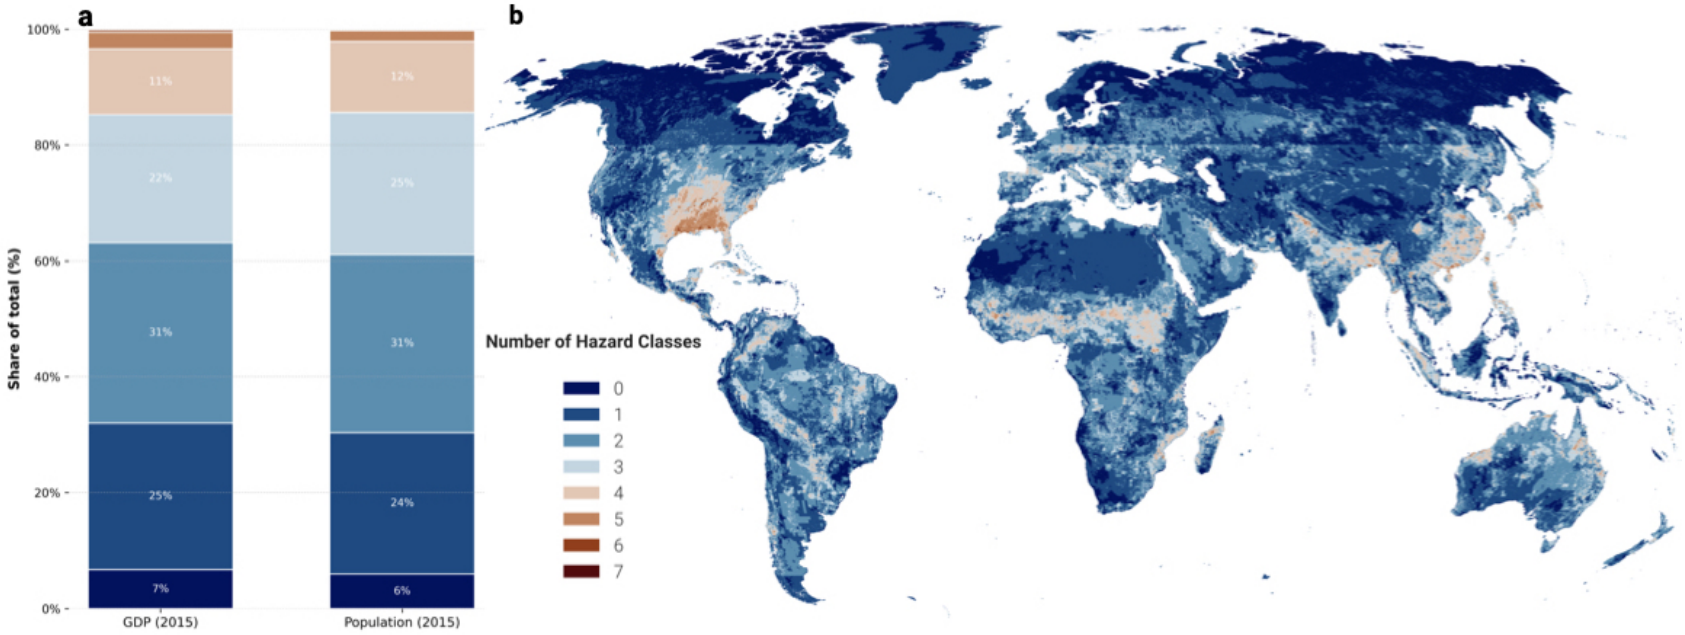

Note: This figure is based on the dataset compiled in Reinhardt (2024a) on a subset of eleven common hazard types (drought, heatwave, flood, hail, earthquakes, tropical storms, extreme winds, landslides, tsunamis, volcanoes and wildfires). The figure shows the number of distinct hazard classes recorded between 2004 and 2015. Panel (a) additionally computes the share of global GDP (left) and global population located in areas with the respective number of distinct hazard classes. Both, population and GDP, are computed based on the respective 2015 values. GDP data is drawn from Kummu, Kosonen, and Masoumzadeh Sayyar (2025). Panel (b) breaks down the global estimate into regional subestimates, showing the share of GDP and population by distinct hazard class.

disentangling of the different shocks. As mentioned in Botzen, Deschenes, and Sanders (2019), many empirical studies rely on the EM-DAT database (CRED and UC Louvain, 2024), which provides information on natural disasters globally. The data is widely used, yet it does not report the precise administrative boundary for all events, making it difficult to connect it to regional and local economic studies. Furthermore, criticisms of this data include its tendency to overestimate the wealth of potentially wealthier countries, due to its emphasis on economic damage over hazard intensity (Felbermayr and Gröschl, 2014). Although numerous datasets have been proposed in recent years, it is only with the recent release of the MYRIAD-HES database (Claassen et al., 2023a) that a comprehensive dataset for local hazard intensities has become accessible. This data is compiled by defining 11 event types as hazardous events by setting thresholds and then compiling the data for the years 2004 to 2017 on a daily level. Chapter 2 of this thesis suggests a new dataset since, during its development, there was no standardized dataset that reliably included hazard intensities. The Platform for Economic Analysis of Climate Hazards (PEACH) provides a comparable dataset with slightly different input data on a global scale but on a monthly resolution. Whenever possible, it refrains from applying any thresholds, allowing researchers to use monthly intensities to compute damage and response functions. Unlike EM-DAT, the goal is not to provide estimates of losses, but to enable researchers to calculate for themselves the response of a chosen asset class (such as agricultural output, population, or building infrastructure) to the intensity of a hazard event.

Besides practical concerns, there are multiple empirical reasons why more papers do not include a multi-hazard perspective. First, the same standard econometric issues apply when using compound hazards as for any climatic event. This includes the non-linearity of effects, the temporal and spatial delay of effects, as well as challenges in designing the correct error term structure (Hsiang, 2016). However, another challenge is how to properly account for the interaction of different types of hazards, environmental factors, or climate impact drivers. A good starting point might be to consider interactions between the intensities of two events. However, while this is relatively straightforward with two factors, it becomes increasingly challenging when additional event types are added. With two intensities, for example, temperature and rainfall, and one outcome, one obtains two hazard-specific responses and one high-dimensional interaction term. Figure 3 visualize this using an exemplary visualization studying the effects of temperature and rainfall on plant health in Nigeria. Plant health is measured by the Normalized Difference Vegetation Index (NDVI). The middle panel displays the

independent effects of temperature and rainfall on plant health, while the right panel illustrates the interactions between them. The figure demonstrates that accounting for interactions changes the estimated effects of the individual factors. If a third component were added, such as wind speed, it would become challenging for researchers to determine which combination drives the observed changes. A remedy might be using an analysis of the sensitivity of a specific system, like a "stress test", where it is analyzed at a specific moment and at a specific location, which factors contribute the most to changes in desired outcomes (cp. Zscheischler, Martius, et al., 2020).

Second, different hazard classes might be correlated either through a shared underlying cause (in the case of teleconnections such as the El Niño-Southern Oscillation (ENSO) or Madden-Julian Oscillation (MJO)) or directly physically related like in the case of air pollution and wildfire (see e.g. Burke, Driscoll, et al., 2021). If the correlated hazard classes are excluded from econometric models, they might cause an omitted variable bias. In addition, as mentioned by Hsiang (2016), these unobservables are likely to be autocorrelated over time, which causes standard errors of the regression to be biased. Even if included, they might limit the possibility of precise inference in parametric regression models, potentially due to issues such as imperfect multicollinearity or inadequate modeling of their joint effects. In the scientific literature, in such cases of variable dependence, statistical methods, such as Copula functions, have been used to better account for the interdependence of variables (Zscheischler, Martius, et al., 2020).

Lastly, there is the practical concern of finding the right econometric model. Standard parametric models often fail to account for the complex, non-linear interactions between multiple hazards as those shown in Figure 3. Extreme events, often treated as singular natural experiments, require the use of causal models to identify effects robustly. Regression discontinuity (RD) designs, particularly in spatial contexts, often fail to account for unobserved spatial path dependence at the local level. Bordering regions usually differ in their exposure to climate risk, economic growth trajectories, and climate adaptation strategies, complicating the use of spatial RD designs. Randomized control trials are not applicable given that the treatment of extreme events may be random in time but not in space. Synthetic control methods face challenges in the context of extreme events due to potential general equilibrium effects and spillovers that may bias estimates even with carefully constructed control groups. Finally, and this method has been used extensively throughout the thesis, there are difference-in-difference models. These models allow capturing the effect by comparing exposed areas with never-treated or not-yet-treated areas. A key challenge is ensuring the va-

lidity of parallel trends assumptions, particularly when multiple hazards are present or when extreme events reappear during the study period. To account for these complexities, all chapters of the thesis rely in some form on recent estimators from the difference-in-difference literature. Although not perfect, some of these new estimators enable treatment reversal, which is crucial in disaster studies.

Fig. 3. Example for Non-Linearities in Impact Dimension

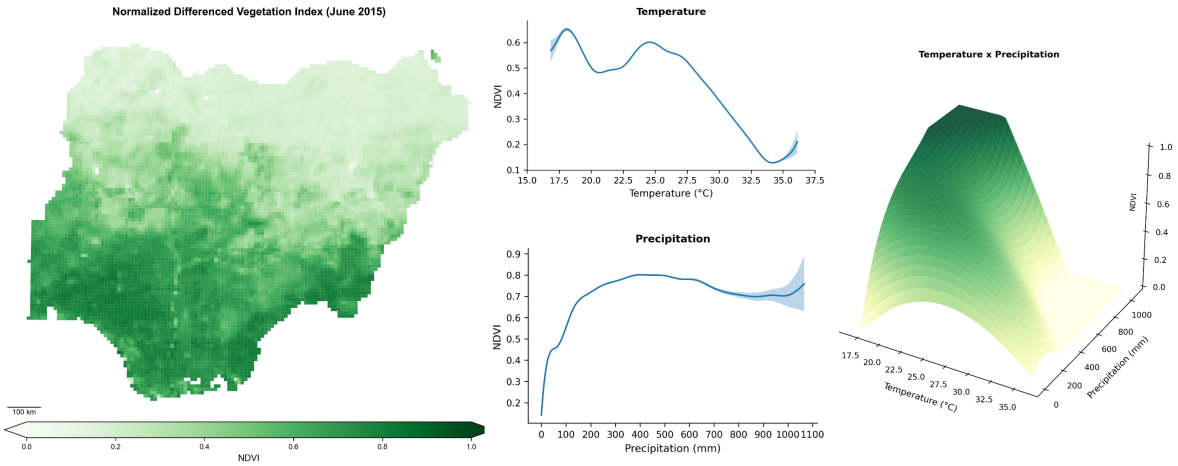

Note: This figure provides an example on the effects of temperature and precipitation on plant health, measured by the Normalized Differenced Vegetation Index (NDVI), in Nigeria. The left figure illustrates, for illustrative purposes, the NDVI distribution in Nigeria for one random month, June 2015. NDVI data is derived from Li, Cao, Zhu, et al. (2023). The two figures in the middle plot a response function of average monthly temperature and mean precipitation in mm onto NDVI by using a generalized additive model (GAM) with four polynomials. Data is derived from ERA-5 (Muñoz Sabater, 2019) and CHIRPS (Funk, Peterson, Landsfeld, Pedreros, Verdin, Shukla, et al., 2015). The right plot shows the interaction of both coefficients on NDVI. The derived data is based on Chapter 2 of this thesis.

In the following paragraph, it is described how this thesis defines connected events and how it addresses the methodological concerns mentioned above. Chapters 2 to 4 include the notion of compound or connected events (called "complex" events in Chapter 3). This thesis adopts a simplified approach by defining a compound event as an event that occurs in the same area ("spatial compound event") or within a certain time period ("temporal cumulative event"). In Chapter 4, also explore different combinations of sequences of events are explored.

To account for interdependencies between hazard types, Chapters 3 and 4 rely on interaction terms. Although continuous hazard measures would have allowed for the estimation of the response depending on a hazards intensity (so-called dose-response

functions), unfortunately such intensities as described in the example of Figure 3 do not reflect the majority of the variables available for this thesis. Notably, due to substantial heterogeneity in the formats of the data used another approach is chosen. Some hazards studied are reported categorically, others through normalized scores, and some in binary form. This inconsistency between data formats substantially complicates the construction of interaction pairs and limits the suitability of non-linear models. Hazards, such as flooding, for example, are more difficult to fit into a non-linear framework. Thus, instead of relying on continuous intensity measures, individual events are transformed into binary treatment dummies based on statistical ("upper tail events") or literature-based thresholds. For some hazards, such as temperature anomalies, upper tail events are defined relative to location-specific historical trends. For other hazards, such as tropical cyclones, a global threshold is used. By using binary treatment dummies this dissertation closely follows the literature on event studies, which is typically based more on binary treatment dummies. To allow for some form of nonlinearity, Chapters 3 and 4 include and discuss quadratic terms in the baseline parametric models. It has to be acknowledged that this comes at the cost of missing parts of the underlying heterogeneity in the impact of events, which is discussed in more depth in the corresponding chapters. Consequently, Chapters 3 and 4 also highlight other factors that researchers may want to consider when assessing whether their findings are truly causal. Another contribution of this thesis is the introduction of a novel dataset, PEACH, in Chapter 2, which addresses the described problem of the lack of a unified database for disaster analytics. Concerns related to temporal and spatial correlation, as well as the error term structure, are addressed by applying different levels of temporal and spatial aggregation. All models from Chapters 2 through 4 incorporate the notion of CID, but call them in Chapters 3 and 4 "extreme events" or weather shocks for readability. Chapter 1 offers an alternative perspective on the literature. By investigating investment flows, the focus of this chapter is on the argument that the effects of destructive extreme events go beyond the sole destruction of physical capital. Regions that are spatially or economically connected with an affected area might be indirectly affected by the shock, especially when shocks occur upstream, so the supply-oriented side of the firm. Although multi-hazard risk is not the primary focus of the article, the model's baseline still accounts for the presence of other extreme events.

The contribution of the thesis is therefore threefold: Chapter 1 highlights, through the example of foreign investment, that the effect of extreme events goes beyond the

sole destruction of assets and might ripple through supply chains into distant firms. Chapter 2 presents a new open-access dataset to better estimate the interactions of extreme events on a local scale. Lastly, Chapters 3 and 4 revisit widely researched topics in development economics, focusing on the interactions between different types of hazards and their implications for the results.

Detailed Contribution of Each Chapter of the Thesis

In the following paragraphs, each chapter is introduced and the respective chapter is summarized, which contributes to the literature in the context of extreme events and connected events in the Global South. The main research question of this dissertation centers around the effects of extreme events, interacting or not, beyond pure destruction, on economic outcomes in countries with little to no resilience to these shocks. Each chapter addresses this question in some form with the latter shifting more discussion on the "compound" events.

Chapter 1

The first chapter is the only chapter that does not particularly focus on connected shocks or cumulative shocks. Instead, the paper examines the impact of destructive extreme events, particularly earthquakes, on foreign investment flows within a country. Although this question has been substantially addressed at the country level, there are only a few studies that focus on evaluation within the country. Indonesia is selected as the study country because it experiences a high frequency of various extreme events. The article combines data on Foreign Direct Investment (FDI) inflows from 416 Indonesian districts with data on 896 geophysical earthquake event footprints from the United States Geological Survey (USGS) ShakeMap database. In contrast to previous studies, it applies a condition that the labor market needs to be, to some extent, disrupted and physical damages are observed. To consider both spatial and economic links between regions, the dataset is enhanced with a 2016 Multi-Regional Input-Output Table (MRIO) for Indonesian provinces. This allows studying the transmission of effects through domestic supply chains. Applying an event-study estimator, the study finds that investment flows are temporarily reduced by almost 90% shortly after an earthquake. The effects are particularly found in areas with lower historical risk exposure and higher FDI stock, suggesting that risk anticipation and invested sum matters for

the response to a shock. However, the shock is not exclusively driven by exposure of the sector-area itself but also by shocks occurring notably on the supplier side of the foreign firm. Here, the MRIO is used to study the transmission of the shock through investment networks. The study is one of the first that documents these effects on a local level without requiring firm-level data. Although not the main focus of the article, it accounts for the presence of potentially compounding events by excluding other types of hazards, such as floods and wildfires, from the baseline regression. Overall, the study highlights on the case of a single hazard type and a single country that, especially for the analysis of hazard impacts on a local scale, it is also important to consider indirect effects such as spatial or network spillovers.

Chapter 2

The second chapter attempts to bridge the gap mentioned above by providing a global gridded compiled and homogenized dataset for 14 climate impact drivers and 3 non-climatic impact drivers. The dataset combines a large set of climatological variables for each month from 2000 to 2019 with full coverage of the 17 hazard classes from 2004 to 2015. The dataset is compiled on a 10x10 km grid and provides the intensity of an event for the majority of hazard types. In contrast to previous studies, for the most important dimensions, such as drought, heatwaves, temperature, rainfall, and air pollution, the data compilation refrains from imposing a threshold, whenever possible, to allow researchers to choose to set their own respective intensity. This shall enable better calibration of damage functions of economic outcomes depending on the hazard intensity. The dataset helps to better understand the global distribution of extreme events over time and especially facilitates researchers in studying the impact of connected or simultaneous events on economic outcomes. To ease usage, it also includes a set of precompiled socioeconomic indicators for each grid cell, such as nighttime light, population, infrastructure information, and land use. This chapter thus contributes to the literature by providing researchers with an additional easy-to-use dataset to study, complementing preexisting datasets. The future plan includes offering the complete executable replication code, allowing researchers to incorporate additional hazard categories not previously captured.

Chapter 3

The third chapter incorporates the novel dataset from Chapter 2 and applies it to the context of population change and migration in Sub-Saharan Africa. Most of the literature focuses on the effects of specific hazards on households' decisions to migrate. In contrast, this article selects nine fairly common hazards in Sub-Saharan Africa (temperature anomalies, rainfall shocks, heatwaves, droughts, storms, flood, hail, wildfires and locusts) and combines them with estimates on net population change between 2000 and 2019. The chapter and the used panel model study the effects of extreme events on 3863 districts in Africa in two parts: first, it presents results using a more standard approach to investigating the impact, notably by employing binary dummies indicating whether a district is exposed to a specific hazard or not. Second, it is highlighted in the remaining part of the article why the use of binary dummies might fall short in capturing the impact dimension. Instead, the second part of the article introduces a distinction between single hazard events, cumulative events (having two or more events of the same type within a year) and compound events (the overlap of two different event types in the same month). Many of the findings in the first part are actually driven by combinations of events rather than by single events. Based on this finding, three key takeaways are drawn: First, the effects found in standard models may not fully capture the real drivers of population change. Instead, each hazard-pair combination should be studied separately as the impact these combinations have on socio-economic outcomes, such as crop production, food security or income, might vary depending on the considered type of event and pair. Second, most of the effects that induce net population changes are found in poorer, remote and climate-vulnerable regions of Sub-Saharan Africa. The discussed mechanism highlights that the impact on specific important crop production and food insecurity are likely drivers for the effects. Hereby, especially drought and heat-related types of events seem to matter. Third and lastly, the article is careful to refrain from calling the effects causal. In one part, due to data limitations, but more importantly, because robustness checks reveal hidden spatial spillovers and non-linearities. These findings reveal a deeper structural complexity of how extreme events affect socioeconomic outcomes. Missing to test and account for non-linearities, spatial spillovers, dynamic carry-over of effects and interactions of hazard events risks to claim preemptively causal effects that are in reality driven by not captured model dynamics. Therefore, this chapter of the thesis contributes to the literature by actively accounting for within-year hazard interaction types, revisiting potential mechanisms that may explain why extreme events influence population

changes and by highlighting risks that researchers should consider when relying on hazard dummies.

Chapter 4

The final chapter investigates the short-term effects that four of the most common weather shocks (floods, heatwaves, storms and droughts) have on the physical urban expansion of 5 721 Sub-Saharan cities. In contrast to the literature and in line with the aim of this thesis, shocks are added simultaneously and capture the total amount of shocks in a given year by hazard type. Applying a panel model over 20 years, the article documents a significant reduction in urban expansion of 3 to 9% following flood events. The effects are particularly concentrated in poorer, more rural cities in West Africa. The effects are amplified if, within the month previous to the flood, heat waves were present in the city. While this article considers urban expansion in itself to be neutral, it documents the following side effects: an intensification of urban population, increased informal settlements, and a higher share of the local population being exposed to future flood risk. Furthermore, an increased likelihood of violent unrest following such events is documented. Thus, the reduction of urban expansion is connected with undesirable outcomes. Heatwaves primarily affect their surroundings through their interaction with flood events, but they have a minimal impact on their own. Storms are fairly local events that affect only a small fraction of African cities. Although strong positive growth effects are found for wealthier cities, the small sample size of treated cities requires attention in terms of the generalizability of results. Droughts at the city level barely matter. However, the article finds that, considering the urban catchment area, droughts lead to a significant acceleration of urban growth by approximately 3%. By applying a recent event study estimator, it is further found that while drought appears to accelerate urban expansion, if one takes the history of past droughts into account, the effects reverse and an additional drought actually leads to a reduction in urban growth. From a literature perspective, this chapter adds to existing literature that, similar to Chapter 3, certain combinations of events might drive part of the initially observed effects. The effects of weather shocks on urban expansion depend on whether researchers take into account the treatment history, spatial scope and the interactions of weather shocks with each other. From a policy perspective, this chapter highlights that floods appear to add pressure to poorer cities, which experience larger growth in informal settlements and potentially worsen living conditions in these cities.

Together, these four chapters offer insight into the complexity of the impact dimension of hazards and the underlying mechanisms driving these outcomes in low- and middle-income countries. Hazards do not only destroy houses and cause death. Instead, their effects ripple through supply chains, spill over to other regions and might lead to persistent economic effects. Furthermore, they rarely occur in isolation. Global warming and atmospheric conditions are increasing the interconnectedness between extreme events, which may lead to new, understudied impacts. Although the economic literature has begun to account for this increased complexity, more research is needed to improve evidence-based policy recommendations. This thesis contributes to the emerging economic literature on the complexity of hazards' impacts and their effects on economic outcomes, while only scratching the surface of what remains to be explored.

CHAPTER 1

Shaking Up Foreign Finance: FDI in a Post-Disaster World

This chapter is solely authored[1] [2].

1.1 Introduction

The literature on disaster adaptation, particularly in the context of climate change, has seen a growing body of research. However, little research exists on the exploration of the impact of multinational corporations in the post-disaster landscape. This gap is particularly striking given the significant contributions of foreign firms to economic development and their key position in domestic supply chains. It also provides motivation to analyze how investment flows respond in the aftermath of a disaster.

This paper aims to assess the immediate changes in inward Foreign Direct Investment (FDI) flows within a country on a local level, as well as the indirect effects through sectoral supply chain dynamics within the country following an earthquake event. Indonesia provides a good case study to study this for three reasons: First, in terms of population, Indonesia is one of the largest economies in the world that sees substantial economic growth. Consequently, it has increasingly attracted foreign investment. These FDI inflows have been found to contribute to the increase of the productivity of domestic firms (Suyanto, Salim, and Bloch, 2009), to accelerate structural transformation (Steenbergen et al., 2020) and to transition Indonesia to a low carbon economy (Zhu et al., 2016; OECD, 2020). These positive effects on the economy motivate us to use Indonesia as a promising case study. Second, Indonesia is, unfortunately, highly exposed to a wide variety of hazards. In 2019 alone, Indonesia faced 48 natural disas-

[1] I wish to thank Lisa Chauvet, Sandra Poncet, Jérémie Gignoux, Ariell Reshef, Myriam Ramzy, Rémi Bazillier, Clément Bosquet, Andrea Cinque, Fabio Ascione, Thibault Lemaire, as well as the participants of the Doctorissimes Conference 2022, the Joint-Research Workshop between FEPS-Paris 1 in Cairo, the International Conference on Development Economics, the German Development Conference and participants of the Internal Development Seminar at Paris 1 and internal seminar at the University of the United Nations.

[2] This chapter is adapted and extended from the original published research article by Reinhardt (2024c), published by Springer Nature. Reproduced with permission from Springer Nature. The Version of Record is available at: DOI: 10.1007/s41885-024-00148-2.

ters that affected approximately two million people, resulting in around 600 casualties, and causing estimated 1.3 billion dollars in damage based on data from the Emergency Events Database (EM-DAT) (CRED and UC Louvain, 2024). Third, the Indonesian Investment Board (BKPM) provides detailed information on the total inflow of foreign investment at the district level (*kabupaten/kota*). Using district-level FDI information enables a better understanding of the effects that extreme events have on local investment. Considering the significant role of FDI in Indonesia's sustainable development, it is essential to explore and unravel how FDI reacts to the growing challenge posed by extreme events.

To investigate the response of Foreign Direct Investment (FDI) to disaster shocks, we use FDI inflow data spanning 2003 to 2019, covering 416 Indonesian districts. This data is sourced from the Indonesian Investment Board. We link this district-wise data with event-specific geophysical earthquake intensity records for all significant earthquakes that occurred during the same period using the speed of ground movement in a metric called *modified Mercalli intensity*. The data is drawn from 896 earthquake events recorded by the United States Geological Survey. Then, an event-study type difference-in-difference model is applied to account for the temporary dynamic response of FDI inflows. To analyze the indirect effects of a disaster through supply chain linkages, this study adapts a panel model from Acemoglu, Akcigit, and Kerr (2016) to the context of a multiregional input-output matrix over 32 Indonesian provinces. In contrast to previous work, the study is structured to distinguish between direct and indirect effects of a disaster following Antonioli, Marzucchi, and Modica (2022). Direct effects impact a firm's production capabilities, while indirect effects disrupt its relationships with partners and business associates. This distinction is important because neglecting the indirect effects of a disaster could worsen the adverse impact of a disaster shock and thereby underestimate the observed economic effects of a disaster.

When we study the direct effects of an earthquake, which is recorded to be destructive, we observe that affected districts see a reduction of FDI inflows by around 90% in the first year following a shock before returning to pre-disaster levels and remaining thereafter. This decrease in FDI is primarily observed in regions with low pre-existing disaster risk, where such shocks are less anticipated. In contrast, areas with high disaster risk do not experience significant fluctuations in FDI inflows. This underscores the importance of disaster preparedness and anticipation in shaping the response to earthquakes. The indirect effects of an earthquake considered in our study are spatial spillover effects and, notably, shocks propagating along the national supply chain.

Local spillover effects from adjacent districts are insufficiently explained by the effects found, likely due to the high concentration of FDI around economic centers and special economic zones. However, the adverse effects on FDI can be partly attributed to shocks occurring in upstream industries, with the most pronounced impact seen in the manufacturing sector, which is the primary recipient of FDI in Indonesia. These negative effects are more significant if the affected upstream industry is located within the same province. In contrast, demand-side shocks can potentially boost FDI inflows, especially in the construction sector. However, this is typically the case when foreign firms in the construction sector remain unaffected and only their customers are impacted. In the robustness section, we find that similar effects occur for flood events, but FDI inflows recover already within the treatment year.

The paper aims to connect several related literature streams. In particular, the paper contributes to the literature on the impact of disasters on FDI. The vast majority of these studies focus on a macroeconomic dimension and study different samples: on a global scale (for instance Gu and Hale, 2023; Doytch, 2019; Neise, Sohns, et al., 2022), on middle and low-income countries (Drabo, 2021) and on Sub-Saharan Africa (Katoka, 2021). Most papers find a negative effect of disasters on FDI inflows. Only a few papers consider a singular country of study (Anuchitworawong and Thampanishvong, 2015; Friedt and Toner-Rodgers, 2022). Friedt and Toner-Rodgers (2022) focus in their work in-depth on the within-country effects of disasters on FDI. Analyzing the impact of six extreme flood and storm events over 16 of India's provinces, they show that FDI tends to spill over to developed regions with large market potential following the months of an extreme event. Instead of analyzing floods, we focus on earthquakes given their short but highly destructive occurrence, which provides a better cut-off in time than flood events. Unlike earthquakes, floods are more likely to lead to secondary disasters such as the spread of waterborne diseases, potentially inducing bias in post-treatment outcomes. Our study focuses on a smaller disaggregated level, which accounts better for the local variation of earthquake intensity. Lastly, we also consider the sectoral dimension of FDI inflows. In contrast to Friedt and Toner-Rodgers (2022) we do not find long-term effects of FDI reduction, which is likely due to the aforementioned differences in disaster category and level of observation considered.

Studies focusing on an even more disaggregated level utilize firm-level data. Notable examples of that literature are, for instance, a study on the impact of typhoons on plant sales (Elliott et al., 2019), several firm-level studies on the Kobe and Tohoku earthquakes in Japan (Cole et al., 2019; Inoue and Todo, 2019; Boehm, Flaaen, and

Pandalai-Nayar, 2019; Carvalho et al., 2021) or studies of floods on firms in Tanzania (Rentschler, Kim, et al., 2021) or tropical storms in India (Pelli et al., 2023). Relevant for this study is the paper by Brucal and Mathews (2021). They combine data from 1990 to 2015 from Indonesian plants and connect them with the EM-DAT database, on a district level, for flood exposure and find a persistent decline in output per worker for foreign-owned firms. We extend their approach by using a more geolocated version of disaster exposure while putting more spotlight on the sector-network aspects of financial flows.

The above-mentioned studies have improved our understanding of how FDI flows patterns evolve in the aftermath of natural disasters on a macro-level and how firms in general respond to disasters on a micro-level. Our study attempts to build a bridge between these two strands of literature. It connects the within-country analysis of FDI inflow data as in Friedt and Toner-Rodgers (2022), similarly focusing on temporal and spatial aspects of FDI inflow on a smaller geographic unit (district level), with the sectoral-network perspective often found in literature on effects on firms (e.g. Carvalho et al., 2021; Inoue and Todo, 2019; Boehm, Flaaen, and Pandalai-Nayar, 2019; Huang, Malik, et al., 2022).

The rest of this article is structured as follows: section 1.2 presents more context on the FDI and disaster environment in Indonesia along with the corresponding data used in this paper. section 1.3 shows the empirical strategy. Then, in section 1.4 the results are presented divided into direct and indirect effects of earthquake exposure, followed by a brief discussion and conclusion in section 1.6.

1.2 Context and Data

1.2.1 Defining the Earthquake Shock

To study the effects that disasters have on FDI inflows in Indonesia, we particularly chose earthquakes as a disaster category. Indonesia, situated between multiple tectonic plates, including the Australian and Eurasian plates, is one of the world's most seismically active regions. It received tragic notoriety due to the 2004 tsunami, which followed a tectonic shift and killed thousands of people in Indonesia and neighboring nations. Its equatorial position also makes Indonesia susceptible to heavy rainfall, leading to hydrological disasters like landslides and floods. Finally, the northern and

southern regions are occasionally affected by storms and cyclones. However, as shown in Figure 1.1, earthquakes and their associated tsunamis account for more than half of all recorded damage and are the deadliest disaster categories. Floods, in contrast, tend to affect the daily lives of Indonesians more. Despite measures for disaster preparedness through national agencies, disaster risk remains a major operational risk for foreign firms. As earthquakes are the most destructive and lethal disaster category, we chose them as the treatment category[3].

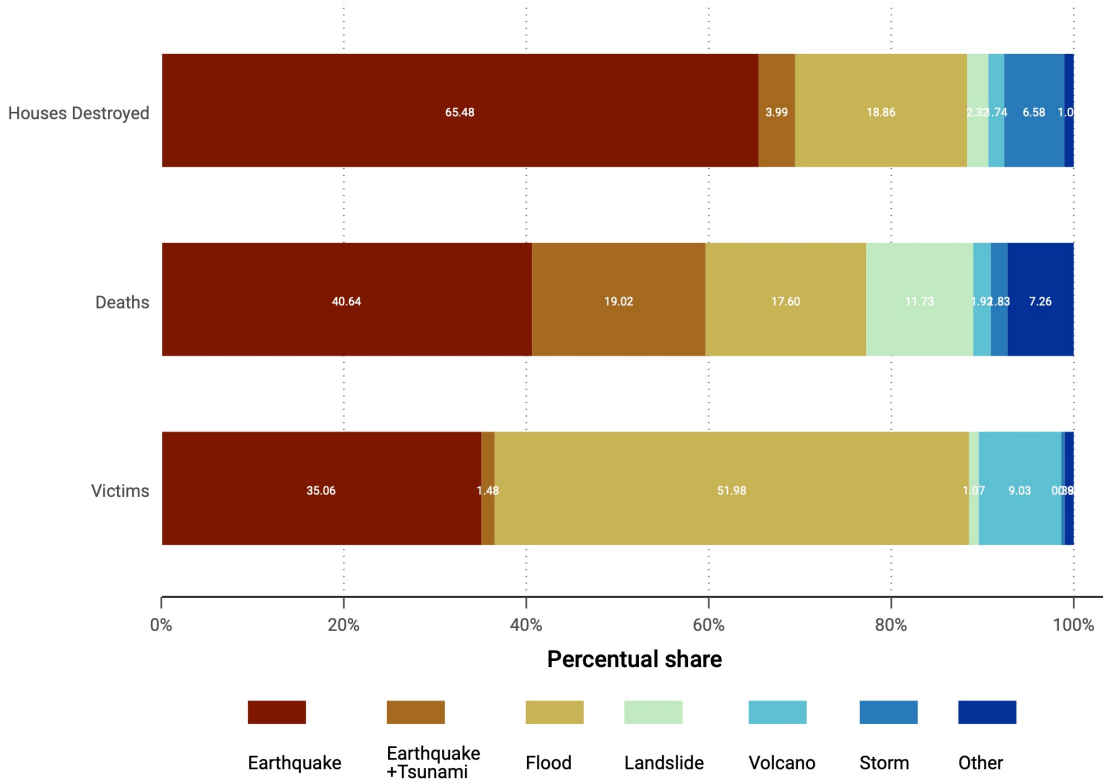

Fig. 1.1. Overview on Indonesia's Hazard Composition

Note: Data is taken from UNDRR (2022b) and aggregated over the period between 2003 and 2019. Calculations are the author's own. The category "other" includes forest fires, tsunamis, droughts, and tidal waves. In absolute numbers there are 10 014 145 victims, 22 454 deaths, and 857 275 houses destroyed recorded. The 2004 Bandah Aceh Tsunami is excluded.

Earthquakes are frequent in Indonesia and affect a large part of the country. To connect them to other variables at the yearly district level, data from several sources are combined. Recent literature has pointed out the issues associated with relying solely on

[3] Yet flood, as the second most common category, is also assessed for robustness.

victim or damage data in disaster studies (Felbermayr and Gröschl, 2014; Felbermayr, Gröschl, et al., 2022). Yet, this information might still be valuable to assess the severity of a shock, especially in local labor markets. Thus, we combine data on the share of the local population affected with an exogenous earthquake intensity measure called Modified Mercalli Intensity (MMI).

The data on the affected population is taken from the Desinventar database, hosted by the UNDRR (2022b) and is, in case of Indonesia, updated by the Indonesian National Board for Disaster Management (BNPB). Similar to the commonly used EM-DAT database, it also records information on damages and the affected population. In contrast to EM-DAT, however, it records information not only on an event-level but also on a geographically narrow area. We draw from Desinventar three variables to measure whether a person is affected: deaths, victims and affected[4]. We build the sum over the three variables given that these people are either permanently affected, by being dead, or indirectly affected, by suffering losses. In both cases, individuals drop out of the labor force. Additionally, we consider earthquake-induced tsunamis victims as with the database, they can be specifically attributed to a related earthquake event. To obtain the share of people affected, the number of people is weighted by the one-year lagged local population with data from the INDO-DAPOER database (The World Bank, 2022b)[5].

The exogenous earthquake intensity within a year and district is based on 897 so-called ShakeMaps from the US Geological Survey (2020) between 6° and -11° geographic latitude and between 95° to 142° longitude. These event-specific maps show the intensity of ground movement in proximity to an epicenter[6]. For each year, we compute the highest measured *modified Mercalli intensity* (MMI) recorded by any of these events within a self-designed 5x5km grid of Indonesia[7]. An example of a singular event in 2005 is shown in Figure 1.2. The MMI intensity scale ranges from zero to ten, while damages on buildings start to occur above level five (United States Geological Sur-

[4] Per UNDRR (2022a) definition direct refers to *victims*:"*persons whose goods and/or individual or collective services have suffered serious damage, directly associated with the event*" and indirect refers to *affected*: "*people, distinct from victims, who suffer the impact of secondary effects of disasters for such reasons as deficiencies in public services, commerce, work*"

[5] The population data is only available until 2014. All years after 2014 are imputed using the district-specific linear population trends

[6] In contrast to the scale, which is measured at the epicenter.

[7] Precisely, the ShakeMap values assigned are values for *peak ground acceleration*, another metric for velocity of the ground, which in a second step is transformed to the MMI based on values provided by Wald et al. (1999). This conversion is also performed by Gignoux and Menéndez (2016)

vey (USGS), 2022). As earthquakes are highly local events, even on a district level, they might not necessarily occur where firms are located. Thus, to improve the quality of the measure, we identify economic centers within each district using nighttime light data by Chen, Yu, et al. (2021) and population density for four time periods by SEDAC-CIESIN (2017) for poorer areas. This is done under the assumption that FDI is allocated, where economic activity is located. We define an economic center as the brightest quintile, where it is possible, of 5x5 km grid cells within each district and year. If the area is small and primarily populated by cities, this method will only identify areas with the most industries. If the area is larger, it will only find the parts that are likely to be populated with cities[8]. Defining the measure on a yearly basis allows for accounting for increasing urbanization tendencies. In some rural regions, there is a lack of sufficient nighttime illumination, which is why the highest quintile of population density is utilized as an alternative measure, assuming that firms operate in the vicinity of the population. The final variable will be the highest MMI value recorded within the economic center of a district within a year. To obtain a binary dummy, only districts with an MMI above five are considered treated as with this intensity minor damage can be expected.

The final treatment dummy combines both datasets in the form of one binary dummy if the local population is affected, the labor shock, and one binary dummy if districts are exposed to a potentially destructive earthquake, the capital shock[9]. Both shocks, labor and capital, can lead individually to a re-evaluation of the foreign firm's investment decision, but together they reflect the complex consequences of a disaster. If the labor market is disrupted, foreign firms with reserve assets might be able to hire workers from farther away, while if capital is affected, firms might substitute, if possible, with a larger work force. However, if both inputs are concerned, foreign firms might not be able to cope so easily. Thus, the final treatment dummy is the product of the

[8] In relative terms, this approach identifies 7% of districts as economic centers. This value reaches 20% in the highest illuminated districts. The value is lower on average due to many non-illuminated cells. As a quintile approach is used the spatial size of an economic center depends on the size of the district. To validate whether nightlight indeed proxies economic centers, it is cross-referenced against city geolocations in subsection A.8

[9] The precise definition for district i and year t is

$$\text{Shock}_{it} = \underbrace{(\text{MMI}_{it}^{\text{urban}})}_{\text{Capital Shock}} \times \underbrace{(\frac{\text{Deaths}_{it} + \text{Victims}_{it} + \text{Affected}_{it}}{\text{Population}_{i,t-1}^{\text{imp}}})}_{\text{Labor Market Shock}}$$

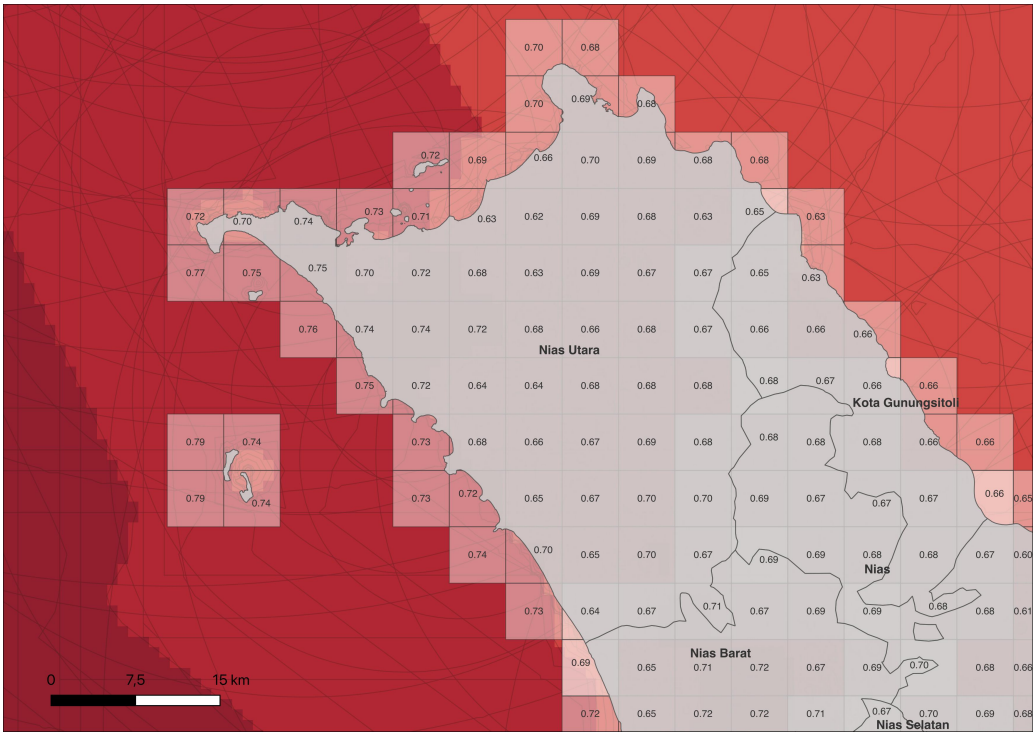

Fig. 1.2. Example of a District Exposure in 2005

Note: The figure shows the exposure of the Nias district to the earthquakes taking place in 2005. It is overlaid with the 5x5 km grid. The red background color indicates the discrete predicted earthquake intensity. A darker red indicates a closer proximity to the epicenter, thus a stronger intensity. The numbers show the mean *peak ground acceleration* within the given grid cell. The map is created in Quantum GIS.

two dummies. In the appendix, shocks individually are also assessed and more information is provided on the chosen threshold of MMI VI. This and more information on the robustness of the shock metric can be found in Appendix A.7.

1.2.2 Foreign Investment into Indonesia

The contribution of FDI inflows to Gross Domestic Product (GDP) in Indonesia mirrors the trends in neighboring Southeast Asian nations. Initially, they rose steadily, reaching 2.8% of GDP in 2014, but then dropped sharply to 0.5% in 2016, before gradually recovering to 2% in 2019 (The World Bank, 2022a). The drop and the slow recovery can be attributed to stalled merger and acquisition (M&A) processes and also to the withdrawal of some multinational companies (UNCTAD, 2017). Jakarta, Indonesia's capital, initially received the bulk of FDI, but its share decreased from 60% in 2003 to approximately 15% in 2019. In contrast, West Java, which is close to Jakarta, experi-

enced a consistent increase, accounting for approximately 20% of all FDI inflows in 2019. Regarding sectors, manufacturing attracted the most FDI, with greenfield investments being prevalent. Agriculture and services saw larger shares of M&A. However, manufacturing's overall share of FDI inflows has declined recently due to difficulties in recruiting workers OECD (2020). These trends are illustrated in Figure A.1. From a geopolitical point of view, between 2010 and 2019, more than 50% of incoming FDI was sourced from ASEAN countries, notably Singapore, followed by Japan (30%), the EU (6%) and China (5%) and Hong Kong (3.5%) (OECD, 2020, Ch.2). Increasing disinvestment has been observed from some EU countries as well as the US. These aggregate flows are not specific to the post-disaster response of foreign investors but highlight the importance of key foreign players in the Indonesian market. Although systematic studies quantifying the concise response of FDI fluctuations after disasters within Indonesia by origin are limited, the strong investment from Singapore, Japan and to some extent China may reflect strategic interests. More research is needed to evaluate whether such interests also shape the post-disaster behavior of foreign investors operating in Indonesia. To study the effect on FDI inflows, the article relies on data from the Indonesian Investment Co-ordinating Board (BKPM) (2021) (*Badan Kordinasi Penanaman Modal*). The data contains extensive information on sectoral investment for foreign investment since the early 1990s and is based on firms holding business licenses, excluding investments through mergers or acquisitions. It is measured in terms of 1000 USD per year. In total, the FDI data extends over 24 sectors and is later assigned an industry code using the ISIC Revision 3 classification. Unfortunately, the database does not capture all levels of FDI received as it excludes investments in oil and gas, banking, financial and portfolio investment, as well as household investments, as it focuses on issued business licenses. The other investment types fall under different jurisdictions, which is why they are excluded from the data set. According to OECD (2020) the data covers investment realizations by foreign capital investment companies, including those below 10% share and with joint ventures with local partners. This practice is likely to inflate the FDI inflow recorded. However, it is the only data available at a second-level administrative level. In order to be able to connect the data later, the investment information by district is geocoded by using the second-level administrative boundaries from OCHA's Regional Office for Asia and the Pacific (2021). All additional sources used are shown in Table A.1.

Table 1.1
Descriptive Statistics

	(1) Full sample	(2) Control	(3) Treated (MMI6+)	(4) Difference (3)-(2)
Dep. Variable				
Arcsin (FDI inflows) in 1000 USD	5.70	5.82	4.87	***
	(5.379)	(5.458)	(4.726)	
Main controls				
Log of mean homogenized NTL in district	-3.32	-3.28	-3.59	**
	(2.897)	(2.910)	(2.788)	
Log of population (imputed after 2014)	12.78	12.81	12.56	***
	(0.928)	(0.923)	(0.933)	
Risk Measures				
Composite Multi-Hazard Index	13.43	13.35	13.97	***
	(2.681)	(2.713)	(2.381)	
Earthquake: Mean std. Risk	2.07	1.95	2.91	***
	(0.856)	(0.786)	(0.854)	
Additional Controls				
Arcsinh (Domestic Inflows) in 1000 USD	5.85	6.08	4.28	***
	(6.560)	(6.615)	(5.936)	
Arcsinh (World Bank Foreign Aid)	14.80	14.87	14.36	***
	(1.682)	(1.664)	(1.745)	
Log (People in Labor Force)	12.06	12.09	11.83	***
	(0.948)	(0.942)	(0.959)	
Poverty Rate (in % of population)	13.75	13.46	15.73	***
	(7.998)	(8.026)	(7.519)	
Morbidity Rate (in %)	29.09	29.10	29.00	
	(8.062)	(8.030)	(8.287)	
Bartik-like Migration Control (standardized)	0.02	0.04	-0.10	***
	(1.052)	(1.118)	(0.326)	
Manufacturing Share (as share of total GDP)	0.15	0.16	0.08	***
	(0.147)	(0.153)	(0.0626)	
Observations	4838	4224	614	4838

Note: Mean coefficients; standard deviation in parentheses. $* \ p < 0.10$, $** \ p < 0.05$, $*** \ p < 0.01$

1.2.3 Discussion

The final dataset is a balanced panel of Indonesia's 514 *districts*[10], the second administrative level, for 17 years, of which 422 are included in the analysis. This is due to two reasons: first, not all districts have sufficient data available, but notably, second, due to

[10] Indonesia is divided into provinces (level 1), regencies and cities (level 2), districts (level 3) and villages (level 4). So the second administrative level refers to 416 regencies (*kabupaten*) and 98 cities (*kota*), but for simplicity all second-level entities are called districts to align the name with the common perception of districts being the level below provinces.

Indonesia's decentralization policy in 2001, many districts have changed their names. (OECD, 2020, chapter 7). To avoid FDI values being assigned to the wrong districts, all observations are standardized to the identifiers of the Indonesian statistical agency (*Badan Pusat Statistik*) in 2003. If a district has been divided after 2003, all offshoots are collapsed into the one district named as in 2003. An overview of the final data is provided in Table 1.1. The treated and control districts are significantly different along the dependent and independent variables. This is not unexpected given that earthquake exposure might, after all, have substantial effects on economic and social development. Although earthquake events are random in time, they may not be in space, which is why economic and human activities tend to self-select accordingly. The resulting selection into treatment of foreign firms will represent one of the key limitations of this paper. In terms of systematic differences between districts, the paper will include differential controls in an extended baseline model.

1.3 Empirical Strategy

The paper investigates the impact of earthquake exposure from three different angles: first, it takes a temporal perspective to see how foreign firms respond to a disaster shock. Second, regarding spatial aspects, the paper examines whether the impact of an earthquake extends to proximate and neighboring districts. Lastly, the paper examines the impact on a sectoral level, but not only within the same industry, but also in linked industries via forward and backward linkages.

1.3.1 Estimation Strategy

The temporal perspective

For the temporal perspective, an event-study type difference-in-differences approach is used of the following form:

$$f_{ijt} = \alpha + \sum_{\substack{i=-4 \\ i\neq -1}}^{4} \gamma_c D_{ij,t+c} + \Gamma X_{ij,t-1} + \omega_i + \omega_t + \omega_{jt} + \epsilon_{ijt} \tag{1.1}$$

The dependent variable f_{ijt} is the hyperbolic sine transformed value of FDI inflows in district i in province j and year t. D_i in Equation 1.1 refers to a dummy, which is based on the computation of the total shock mentioned in section 1.2. The dummy

equals 1 if persons have been recorded as directly affected by an earthquake and the economic center of the district is exposed to an earthquake intensity with at least minor damage potential. Otherwise, if only one or none of the conditions is fulfilled, the dummy equals zero. Considered as treated are only districts with a single earthquake event, given that more than two-thirds of the treated sample have only a single event recorded. Districts with more than one recorded earthquake are excluded from the analysis, which is assessed in the robustness section. In Equation 1.1 also the four leads and four lags around the treatment year of any earthquake between 2003 and 2019 are included to add pretrends and post-treatment periods, while excluding the first lead. These leads and lags are created by using the years surrounding the treatment and by conditioning that only a single earthquake occurred (as treatment reversal complicates the baseline DiD model). In all models $X_{ij,t-1}$ contains the constant, the logged total population and the inverse-hyperbolic sine transformed nighttime-light intensity lagged by one period to account for the population affected and the economic situation of the district. ω are the district, year and province-year fixed effects. They shall rule out unobserved time-invariant heterogeneity within districts like distance to the coastline or availability of a port, time-variant constant shocks across years like the financial crisis shock of 2008 and for instance time-varying province-specific regulatory changes and policies due to province-level political decisions. The standard error is clustered at the district level. To isolate the effect of earthquakes on FDI without contagion through other disasters, a dummy is created to control for simultaneously occurring other disaster types[11]. Every baseline regression is conditioned on the dummy for simultaneous disasters being 0 to ensure that no other disaster than earthquakes materialized within year t.

The spatial perspective

Papers such as Felbermayr, Gröschl, et al. (2022) or Friedt and Toner-Rodgers (2022) use models relying on spatial weighting matrices in their spatial specification. This paper uses a simplified version of their approach by first identifying neighboring districts and then by applying a modified version of Equation 1.1 using an interaction term:

$$f_{ijt} = \alpha + \sum_{\substack{i=-4 \\ i\neq -1}}^{4} \gamma_c D_{ij,t+c} \times \gamma_c D^{k\neq i}_{kj,t+c} + \Gamma X_{ij,t-1} + \omega_i + \omega_t + \omega_{jt} + \epsilon_{ijt} \tag{1.2}$$

[11] Based on victim records within the Desinventar database. This dummy equals one if more than 20 people are affected by any disaster type other than earthquakes or earthquake-induced tsunamis.

In essence, the model compares the changes of FDI inflow if a district *i* is treated with a non-treated district under the condition that either none or at least one adjacent neighboring district *k* is also treated at the same time.

The sector-network perspective

For analyzing the sectoral outcomes, the paper broadly follows the approach by Acemoglu, Akcigit, and Kerr (2016). Using a multi-regional input-output table for Indonesian provinces from 2016 (Badan Pusat Statistik, 2021), up- and downward linkages are calculated[12]. Upstream linkages are used to show the impact of a supply-side shock on foreign investment, whereas downstream linkage shocks correspond to the demand-side shock. The shock is different from the one in the temporal and spatial approach. The methodology still relies on the dummy variable D as in Equation 1.1. Yet two extensions are included: the sectoral dimension is added and the input-output data is at the province-level only, so the shock needs to be aggregated and account for the sector dimension. To incorporate both the new within-sector shock $Shock_{sjt}$ in sector s, province j and year t over district i is defined as

$$Shock_{sjt} = \sum_{i}^{i \in j} w_{ijs} \times D_{ijt} \quad \exists w_{ijs} = \left(\frac{\sum_t GDP_{ijst}}{\sum_t \sum_i^{i \in j} GDP_{ijst}}\right)$$

The sectoral shock decomposes into the sum of two parts: a district-time-invariant share of sector relevance within a province and the previously used sector-invariant but time-variant dummy on earthquake treatment. The district-time-invariant share of sector relevance is the mean GDP within a district and sector over the mean GDP within a province over all 17 years. The interaction of district-sector weight and dummy of exposure is then aggregated at the province level. The total shock corresponds to the share of sector GDP affected within a year and province. For instance, if only one district is affected within a province but it makes up 80% of the production within a year, then the shock variable is 0.8 for this sector and this province within this year.

Then, to account for shocks along the supply chain, the time-variant shock by sector and province is combined with information from the input-output Leontief inverse matrix. The resulting supply-side upward or demand-side downward shocks are created

[12] The I-O table is sector-wise aggregated to match the corresponding FDI sectors and converted to the Leontief inverse matrix

in the following ways:

$$Upstream_{jst} = \sum_{k} (Input^{2016}_{k \to s} - 1_{k=s}) \times Shock_{kjt}$$

and

$$Downstream_{jst} = \sum_{k} (Output^{2016}_{s \to k} - 1_{k=s}) \times Shock_{kjt}$$

Hereby, input refers to row entries of the Leontief inverse matrix, whereas output refers to column entries of the Leontief inverse in each respective province and sector. The 1 is an indicator function if the connected sector is the sector in consideration, so if $s = k$ (to avoid double accounting from within sector shock and again shock through linkage values). In other words, the up- and downstream shocks are linkage-weighted sectoral shocks that occur in other sectors than the sector s and within the same or other provinces in Indonesia. The final model to be estimated has the following form:

$$f_{jst} = \gamma f_{js,t-1} + \beta_1 Shock_{js,t-1} + \beta^{\text{up}} Upstream_{js,t-1} + \beta^{\text{down}} Downstream_{js,t-1} + \Gamma X_{j,t-1} + \epsilon_{jst} \tag{1.3}$$

Here the dependent variable is again the arcsin transformed FDI inflow, yet not within a district but on a province level j and by sector s. Then a lagged variable of the dependent variable based on the model by Acemoglu, Akcigit, and Kerr (2016) is also included, followed by three shocks defined above: the sector-specific direct shock, a shock in the upstream and a shock in the downstream sectors to sector s. The X_{jt-1} encompasses the aggregated mean nighttime light and the population by province, which are logged and lagged by one year. Equation 1.3 compares the effect of within-sector-province shocks to those provinces within the same sector that are not or are less affected. The shock variables are measured in percent, so a unit shock within sector-province or upstream/downstream refers to a 1% increased affected GDP share in either the same sector of the supplier- or the customer-network. It is important to mention that, in contrast to Equation 1.1 and Equation 1.2, this model does not condition the regression of only single treatments. If one takes a network perspective, multiple treatments are by definition likely as somewhere in Indonesia over the time earthquakes take place, which is why multiple treatments are explicitly allowed. As an extension, we later also split the up- and downstream shocks of Equation 1.3 into shocks occurring within a province versus shocks occurring outside a province. This will proxy different outcomes from shocks within tight local production networks versus potentially less damaging shocks within nationwide production networks.

1.3.2 Identification Strategy

Identification in the models comes from the random timing and the intensity of an earthquake event. The key identifying assumption is therefore that conditional on controls, so assuming no pre-trends, the timing and intensity of the shock are exogenous. While earthquakes appear random in time, there might be the pressing matter of selection into treatment and thereby violating the parallel trend assumption on which our approach relies. Table 1.1 shows that control and treatment group are different along a wide range of controls and especially FDI inflows are significantly different. This, in turn, might reflect differences in investors' perception of risk. To alleviate concerns of a violation of parallel trends and selection into treatment three approaches are taken: first, in some setups all controls are also included to account for the differences in covariates. Using the three dimensions of fixed effects, district, time and province-year, it is attempted to rule out large parts of unobserved heterogeneity, such as sector-province-specific FDI restrictions or nationwide shocks through the financial crisis in 2008. Second, for robustness a matching approach following Imai, Kim, and Wang (2023) is implemented, whereby control and treated units are matched based on treatment history and a set of covariates. This aims to validate whether the results of the baseline still hold when changing the control group. Third and last, it is tested if the core results hold under violation of the parallel trend relying on a methodology by Rambachan and Roth (2023).

Next, and as one of the largest concerns in this study, there is the risk of a violation of the Stable Unit Treatment Value Assumption (SUTVA), a key assumption that includes that all potential temporal or spatial spillovers are accounted for. It is a key assumption in the difference-in-difference literature and would, for instance, be violated if outcomes were impacted by the treatment status of other units, such as spatially proximate ones or temporally related ones. With Equation 1.2 to a certain extent, spatial correlation is assessed. Finally, all models mentioned above ignore the staggered adoption design of treatment, which according to Goodman-Bacon (2021) might lead to biased estimates due to negative weighting of estimates. In the robustness section, it is assessed to what extent the negative weighting biases the estimates. The baseline regression furthermore excludes districts that have been treated more than once. While our study does not directly account for the frequency of earthquakes, recurrent exposure to shocks might have fundamentally different effects on outcomes. Firms in such district might respond different to exposure given better adaptation and an anticipation of shocks. While not in the focus of this study, for robustness it is assessed in the

heterogeneity analysis if anticipation of risk plays a role, while econometrically these previously excluded districts are investigated separately. First, in Appendix A.3, treatment reversal is explicitly allowed by relying on an estimator by De Chaisemartin and d'Haultfoeuille (2023). In other words, the treatment sample is extended by the districts that have more than a single event recorded. Yet the effects of rarely treated districts and the potential null effects of frequently treated units might cancel out. Thus, in the robustness section, the effects on districts with multiple treatments are considered individually by excluding single treated districts. Therefore, the above-mentioned estimator by Imai, Kim, and Wang (2023) is used to match districts to districts with similar treatment history. These tests serve the purpose on evaluating whether it is justifiable to exclude districts with multiple treatments from this analysis.
Lastly, to account for the complexity of connected disaster types, in Appendix A.6 it is discussed how classes of natural disasters (such as climatological events as a whole) compare to specific earthquake shocks as a robustness check. Hereby, especially flood, as the second most devastating disaster category in Indonesia, is considered in depth in Appendix A.5.

1.4 Results

1.4.1 The Direct Effects of Earthquake Exposure

The key effect of experiencing an earthquake within the economic center of a district can be seen in Figure 1.3. One can clearly observe the drop in negative hyperbolic sine transformed FDI inflows, especially in the year following a disaster. The three models in Figure 1.3 use different specifications: the first is the unrestricted baseline model (dark blue connected line), which follows the model described in Equation 1.1. The second model (lighter blue with a dashed line) additionally excludes districts with another disaster exposure or multiple earthquakes over time, and also districts in which, in the five years prior to 2003 (the first year of observation), disasters occurred. This is to rule out any carryover effect to the treatment period. Lastly, the final model adds a set of extended controls. An overview of the results can be found in Table A.2. In terms of percentage change, the three models would predict a reduction of around 90% in the year after an earthquake significant on a 1% level[13] in comparison to non-

[13] Conversion follows Bellemare and Wichman (2020) (equation 11) of arcsin-dummy conversion: $exp(\beta) - 1$

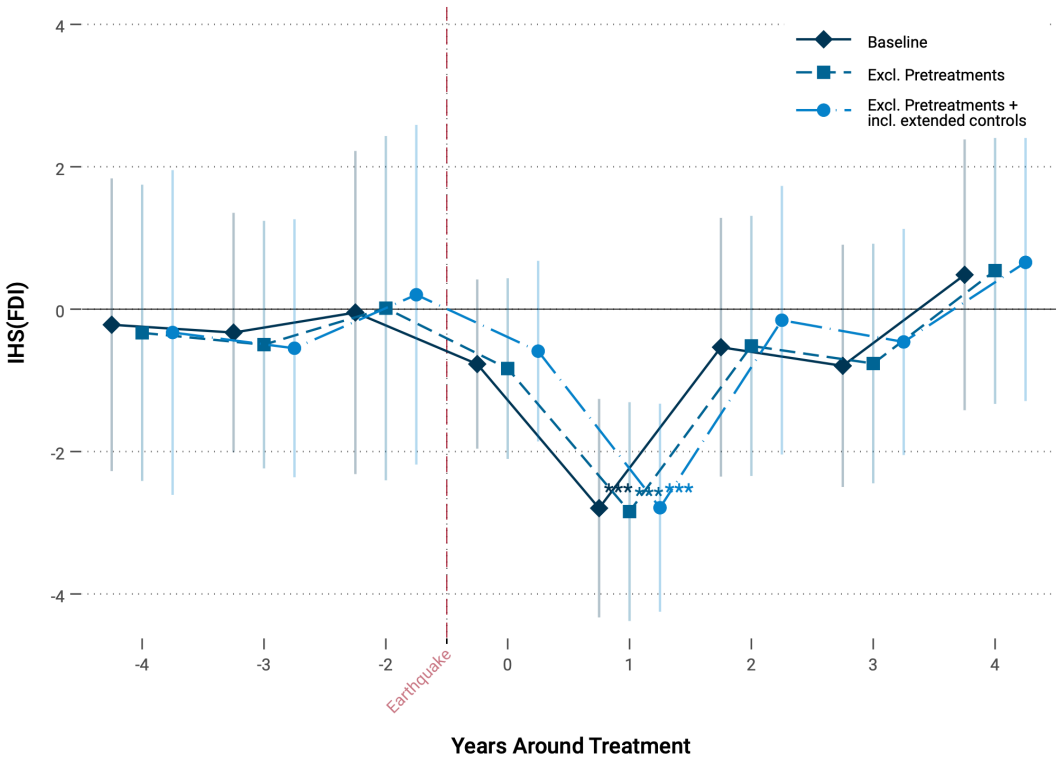

Fig. 1.3. Dynamic Treatment Graph of Exposure to Earthquake

Note: The figure shows the dynamic treatment graph estimated based on Equation 1.1. FDI inflows are inverse hyperbolic sine transformed. The bars depict the 95% intervals and stars indicate *** 1%, ** 5% and * 10% significance. Point estimates are based on at least 4 692 observations with at least 47 districts in the treated group. Each model includes year, district and province-year FE as well as lagged log(nighttime-light) and logged and lagged population as core controls. Models used exclude other simultaneous disaster types and exclude treatment reversal within unit. The corresponding R^2 is approximately 0.68 and the estimates are based on district-clustered error terms..

treated or not-yet treated districts in the same year. For the main controls, only lagged nighttime light emissions are positive and significant. Not surprisingly, the economic proxy predicts higher FDI inflows through more economic activity measured. The log of population density is not significant, likely because aggregated FDI inflows cover all sectors and, for instance, for agriculture population density might play less of a role than for labor-intensive industries. The extended controls in Column (3) show positive effects of domestic investments and a higher manufacturing share within districts. Migration and morbidity are insignificant, whereas a higher poverty rate tends to decrease FDI inflows, which all follow the standard economic theory. Several factors can explain the temporary reduction and subsequent recovery of FDI inflow. Foreign

firms may require time to make decisions on how to respond to the shock, conduct risk assessments for future investments, await improvements in market conditions, or anticipate governmental responses. Once these obstacles are addressed, foreign firms may reinvest. The results do not necessarily conflict with the findings by Friedt and Toner-Rodgers (2022) of a long-term reduction in FDI inflow, as they consider a larger geographic unit and a different disaster type.

On a sector-level, Figure A.2 shows the results of Equation 1.1 if one extends the model to the district level over ISIC Revision 3 classified sectors. On a district level, in the year after the earthquake (right panel) one can observe direct negative effects for all sectors. Manufacturing, utilities and transport, storage and communications tend to be the most affected sectors. The least affected are mining and construction. Nonetheless, all the effects found are comparable to the baseline aggregated estimates.

Robustness Checks

To validate the base results of this paper, multiple robustness checks are conducted to examine the reduction of FDI inflows in the year following an earthquake exposure. First and foremost, there are potential issues with the econometric approach. In the recent Difference-in-Difference literature, the issue of negative weighting is more and more emphasized (see for instance Sun and Abraham (2021) and Xu (2023)). In the current dynamic difference-in-difference estimation this issue is entirely ignored, which might lead to wrong estimates. Figure A.3 shows the same results using recent methods from literature, which account for the negative weighting issues in different ways. The results tend to be comparable to those of the baseline specification in their extent and significance[14]. More information is provided in Appendix A.3.
In the current setup, districts with multiple earthquake events recorded are excluded from the analysis. In order to include them one needs to allow for treatment reversal, to which end an estimator by De Chaisemartin and d'Haultfoeuille (2023) is applied. The results, discussed in Appendix A.3, are broadly comparable.
Another issue arises due to the control and treatment group being significantly different, threatening the parallel trend assumption necessary for causal inference. To reduce the concern, Figure A.4 applies a methodology by Imai, Kim, and Wang (2023). Hereby, treated units are matched with control units with a similar treatment history, then are

[14] Two estimators suggest a small pre-trend. This trend is, however, not significant on any standard level.

matched along covariates and finally differenced using a difference-in-difference approach. Thereby, it reduces the reliance on the parallel trend assumption. Matching is done on a set of socio-economic variables, which aim to mimic investors' risk perception and investment attractiveness of a district. More information on the approach is provided in Appendix A.3. The upper panel of Figure A.4 uses the same definition of treatment as the baseline model of Equation 1.1. All different types of matching or weighting methods show comparable effects of the baseline model, thereby again providing evidence that the core results are not substantially driven by differences in covariates. That the results are robust against linear and non-linear violations of the parallel trends is for robustness also assessed in Appendix A.3 applying a method by Rambachan and Roth (2023). The results support the main findings of this study.
The lower panel of Figure A.4 investigates whether the baseline results vary if one exclusively considers districts with multiple earthquake events instead of those with singular treatments. The same method is applied as for the upper panel. No significant effects are found. This suggests that districts, which are frequently exposed, behave different to those where risk is potentially less anticipated and firms are potentially less adapted. In the next section the role of risk anticipation will be individually regarded. The lower panel provides reasoning that the exclusion of districts with multiple earthquake events seems reasonable given their likelihood of a different response to shocks.
Lastly, it is tested to see if specific outlier events drive the results by dropping one district-event at a time. This does not seem to bias the results.

Heterogeneity : Do Risk Anticipation and Investment Stock Matter?

In this section, the sample of the baseline regression of Equation 1.1 is broken down into a subgroup analysis focusing on risk and investment stock. This is done as there is likely a selection into a treatment pattern: firms settling in earthquake-prone areas might have ex-ante the expectation that earthquakes might emerge as earthquake risk is public knowledge, and start adapting to the disaster risk. Firms in less risky areas, in contrast, might respond strongly to disaster events as their anticipation of risk is lower or nonexistent. Similarly, the response of FDI flows might depend on the stock previously invested. The more one has invested in a region, the higher the individual cost of an earthquake. This subsection investigates the two channels, risk expectation and capital stock.

In order to divide the sample into the relative risk likelihood, we create an earthquake

risk and overall disaster risk (*composite hazard score*) dummy following the idea of Emrich et al. (2022). The thorough computation of the risk metrics and their validation are described in Appendix A.9. Column (1) and (2) of Table 1.2 divide the sample into the above and below median overall disaster risk, while Column (3) and (4) focus on time-invariant earthquake risk. Column (5) and (6) divide the sample into approximate initial FDI stock in 2003. To approximate FDI stock by district, data on FDI stock within Indonesia as a country is obtained for each year from UNCTAD (2021). Then, for each district, the share of individual FDI inflow over total FDI inflow across all 17 years is calculated and interacted with the value of Indonesia's FDI stock in 2003 to obtain a proxy for FDI stock by district. This requires the admittedly strong assumption that FDI inflow is a sufficient proxy for relevance in FDI stock.

The results show that for Column (1) and Column (3) in Table 1.2 significant effects are only found in the treatment year and the year after if the overall risk of disaster or earthquake is below the median national risk. The prediction suggests that FDI inflows are reduced by 90 to 95% being significant on a 5% level in the treatment year and the year after. These effects are not found for above-median-risk districts or after period 1 past disasters. Assuming that disaster risk is public information, these results would align with the hypothesis that foreign inflows are more affected if the shock is less anticipated than in high-risk areas. Focusing on Columns (5) and (6) one can observe that if more initial FDI stock is present in 2003 (Column (6)), one observes a reduction up to 95% of FDI inflows on a 5% significance level in period one after the earthquake, but not for those with lower initial FDI stock. This would confirm the hypothesis that higher initial investment might cause a stronger response to shocks than if less is invested.

1.4.2 The Indirect Effects of Earthquake Exposure

After investigating the repercussions of earthquake events over time and across different subgroups, the indirect impacts on FDI are analyzed, focusing on shock transmission through proximity and network linkages.

Table A.3 presents the impact on FDI inflows based on the interaction terms of district shocks and of shocks materializing in at least one or more adjacent districts. In Column (1) one sees the effect using the definition of shock as mentioned in the data section. One notices a negative outcome of within district shock on FDI of -2.932 significant on a 1% level. If neighbors are also affected, the shock is reduced by approximately

Table 1.2
Dynamic Difference-in-Differences: Heterogeneity

	Composite Hazard		Earthquake Risk		Initial FDI Stock (2003)	
	(1) Below Median	(2) Above Median	(3) Below Median	(4) Above Median	(5) Below Median	(6) Above Median
Event Year	-2.332**	0.171	-2.282*	0.529	-0.475	-1.182
	[-4.382,-0.281]	[-1.836,2.178]	[-4.948,0.383]	[-1.533,2.592]	[-2.230,1.280]	[-5.109,2.745]
Event Year +1	-3.343***	-1.476	-4.505***	-1.280	-1.579	-3.724**
	[-5.744,-0.942]	[-4.215,1.264]	[-7.299,-1.711]	[-3.573,1.012]	[-3.647,0.489]	[-6.616,-0.831]
Event Year +2	-1.359	-0.811	0.201	-0.420	-0.0101	0.142
	[-4.492,1.774]	[-3.663,2.041]	[-2.866,3.268]	[-3.234,2.393]	[-2.533,2.513]	[-6.806,7.090]
Event Year +3	-1.425	-0.820	-1.695	1.013	-0.335	1.299
	[-4.727,1.877]	[-3.456,1.816]	[-5.392,2.002]	[-1.153,3.178]	[-2.277,1.607]	[-3.266,5.865]
Event Year +4	0.00474	-0.699	1.384	-0.0148	1.411	-0.335
	[-2.716,2.726]	[-5.323,3.925]	[-1.157,3.926]	[-2.511,2.482]	[-1.035,3.858]	[-6.014,5.345]
Event Year +5	-2.091	-0.300	0.827	-1.916	-2.482***	0.586
	[-4.940,0.758]	[-2.261,1.662]	[-1.634,3.287]	[-4.234,0.403]	[-4.329,-0.634]	[-2.784,3.956]
N	2679	2005	3411	1353	2265	2472
R^2	0.700	0.715	0.685	0.703	0.545	0.674
Districts Treated	25	23	14	34	32	16
Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	✓	✓	✓	✓	✓	✓
District FE	✓	✓	✓	✓	✓	✓
Province× Year FE	✓	✓	✓	✓	✓	✓

Note: *<0.1, **<0.05 and ***<0.01 p-value. Dependent variable: IHS(FDI inflows). All regressions contain also pre-trends, which are excluded in the output and includes lagged log NTL and log population as controls. All errors are clustered at the district level.

0.5% in the year after the disaster. No other time period is significant. Column (2) and Column (3) change the definition of the treatment dummy. So the treatment dummy in Column (2) considers only districts where persons are recorded to have either died or are directly affected as treated. Using this treatment dummy, in the period after the disaster, all three coefficients are significant in comparison to districts without any shock or shock in proximity. The effect is smaller and predicts with -1.2 a reduction in FDI inflows of only 70%, not like 90% in Column (1). If one adds the neighboring shock, this effect increases significantly at a 10% level. These effects tend to persist for up to two years after the disaster. Column (3) only regards districts where the economic centers are physically exposed to a damaging earthquake. Here, no positive effects are found. As a bottom line, it can be deduced that on a highly local level, it mildly seems to matter if shocks also affect neighbors. The extent, however, remains small and is not entirely clear in terms of direction. The finding that spatial spillovers are common after disasters is also found, for instance, in the growth literature (see e.g. Felbermayr, Gröschl, et al., 2022), but also in the FDI literature (Friedt and Toner-Rodgers, 2022). Yet, in contrast to Friedt and Toner-Rodgers (2022) in this case study no long-term outflow of affected regions is observed and a recovery to pre-disaster trends is found.

Figure 1.4 employs Equation 1.3 and shows how sectoral shocks impact FDI across various sectors at the provincial level through the supply chain. Specifically, it quantifies two distinct effects. First, it examines the sector- and province-specific consequences of an earthquake. In simpler terms, it measures the response of FDI inflows in a particular sector within a province when there is a 1% increase in the affected sector's GDP in that province within a given year. To illustrate, let's consider a scenario where certain districts within a province experience an earthquake during a specific year. If these affected districts contribute to 60% of the province's GDP within a specific sector, the shock variable quantifies how the FDI inflow into that sector changes when one additional percent of the sector's GDP is affected. The effects of higher shock exposure can be found in the first column of the upper panel in Figure 1.4(a) converted into a percent change. Only cells in which the p-value of the coefficient is below the 10% threshold are filled. All other values above are shown in gray. One can notice that especially manufacturing and services appear to see an increase in FDI inflows in the year following the disaster by 7 and 18% correspondingly. The results are different to the ones on a district-level shock, which are strongly negative. Four reasons might explain the difference: first, the shock used is on another scale (one percent GDP shock versus dummy if an economic center within a district is exposed to an earthquake). The

second reason is the difference in aggregation level, on a province level, there might be positive effects, whereas within the district the effects are negative. Thirdly, here disasters are defined on a sector-level not solely on the treatment dimension. Thereby, the treatment group varies, which might also affect the direction of the coefficients. Lastly, the district analysis captures the overall effect on FDI, yet it could also be that FDI inflow is reduced due to supply-chain shocks aggravating the negative effects on a district level, which is discussed in the next paragraph.
The other two columns in Figure 1.4(a) show the consequences of a one percent increased shock through sectoral GDP weighted up- and downward linkage values. In simpler terms, one considers how much a 1% increase in indirect supply and demand side shocks on the change in FDI inflow in the affected sector. The second column upper panel in Figure 1.4(a) indicates that downstream-demand side shocks lead to an increase in FDI received in manufacturing by around 5% and by an astonishing 60% in construction FDI. In contrast, supply-side shocks (last column upper panel) lead only in manufacturing to a reduced FDI inflow of around 13%. Collectively, the tables show that after an earthquake at the provincial level, the manufacturing sector is often the most affected. The shock is not consistently negative; instead, it seems to depend on how the shock affects FDI, with supply-chain disruptions being particularly troublesome for foreign investment. One should keep in mind while interpreting the results as shown in Figure A.1 that due to the used data source the manufacturing sector tends to be the industry with the least non-zero FDI inflow data.

The right panel of Figure 1.4(b) explores another aspect. It looks at both upward and downward linked shocks and distinguishes between two types of linkages: those within the province and those outside the province. The underlying hypothesis is that shocks will have a greater impact within provincial networks, primarily due to the presence of more localized production networks. The first column shows the effect of a within-sector-province shock. A one percent increase in within-sector-province shock increases FDI inflow into manufacturing by 6%, into wholesale, trade and repair by 8% and into other services by 22%. It reduces FDI inflows into construction by 6%. Construction faces a reduction in FDI inflows, likely because foreign firms benefit less from construction contracts if the industry itself is affected.
For a one percent supply-chain shock (next column), FDI inflow into manufacturing is reduced by 16%. A shock in the demand side for goods (downstream) sees an 11% increase in manufacturing and up to 68% in construction. Yet, wholesale, trade and repair FDI inflows reduce up 34%. Service FDI inflows reduce even up to 65%. The

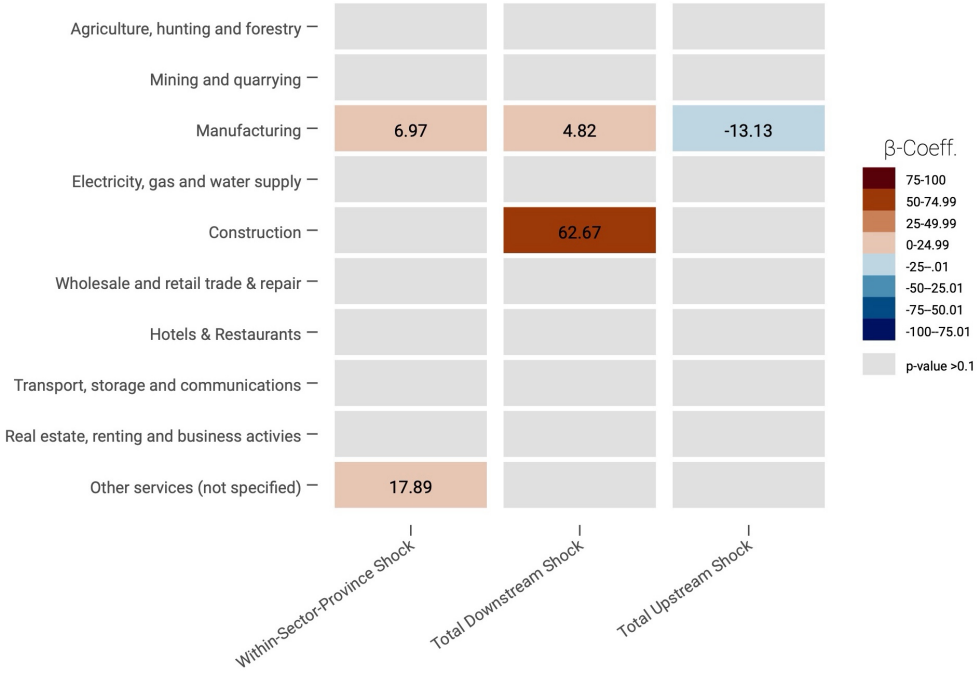

(a) Using Total Linkages

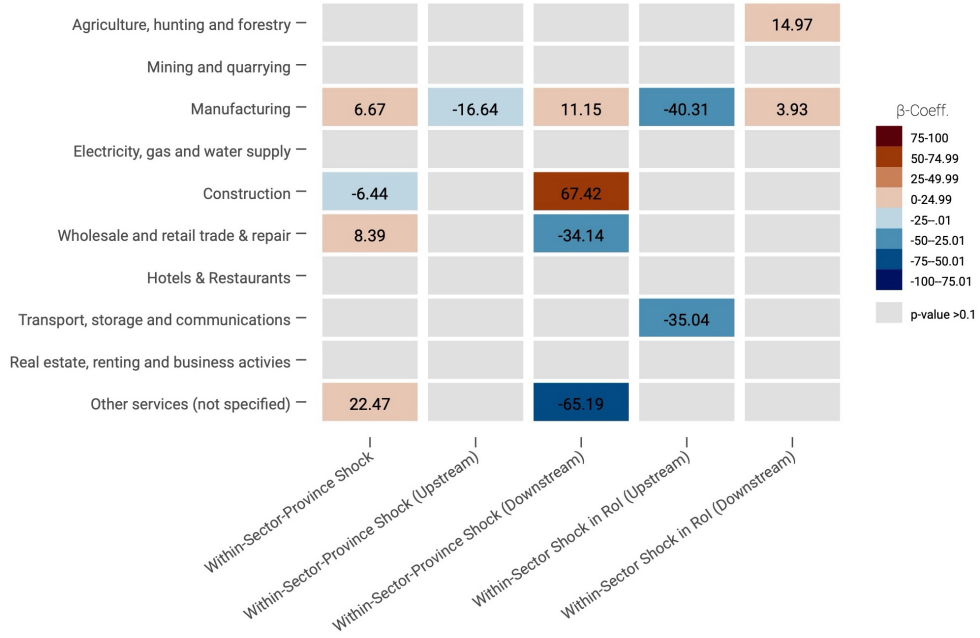

(b) Using Linkage Breakdown

Fig. 1.4. Effect on Sectoral FDI via Linkages

Note: Shown are the β coefficients transformed into percentages. The two panels show the sectoral response of IHS-transformed FDI inflows based on a 1% earthquake-induced GDP shock. The first columns in both panels refer to a 1% shock occurring within the sector of FDI considered. The other two columns to up- and downstream shocks, described as Leontief-inverse matrix weighted GDP shocks. In contrast to the first panel, the second differs also by origin of the shock. Coefficients are transformed using $exp^{\beta} - 1$ as transformation. Grey fields are below the 10% significance threshold. Estimations are based on Equation 1.3. All earthquake events are included and each regression includes lagged and logged nighttime-light and population values as core controls.

final two columns show the effect of a one percent shock within linkages outside the province. A one percent linkage-weighted supply-chain shock in manufacturing reduces FDI inflow by around 40% and transport and storage by 35%. In contrast, agriculture and manufacturing see a rise in FDI inflows concerning a shock on the linkage-weighted demand side.

A one percent increase might weigh more in the regression models for the rest-of-the-world shocks. As mentioned before, the shock here is an interaction term between the province-sector-specific shock and the province-sector-specific linkage value. Table A.4 shows the descriptive statistics of the shocks, and one can notice that the mean of shocks within province-sector linkages are larger within provinces than within the rest of the world, which leads to larger coefficients in the rest of the country metrics as a 1% increase will have a larger weight on the regression.
The found effects illustrate the complexity of effects on sectoral FDI and thereby also aggregate FDI inflows. While specific shocks within the own sector might actually increase FDI received in the year following a disaster, the aggregated effect can still be negative if the production relies heavily on intermediate input goods, which are affected by the disaster, chiefly within the own province. Especially shocks taking place on the customer side tend to benefit the manufacturing and construction sectors, as it is likely that the demand for intermediate goods, for instance, for reconstruction might increase. Shocks occurring upstream reduce FDI inflow as it might cause production interruptions, which is generally bad for business.

Robustness

To validate the effects found in space and sectors, several robustness checks are conducted. First, it is tested to see if one finds reallocation patterns of FDI from affected districts towards proximate economic centers. It is found that this hypothesis does not appear to be true in this study. In contrast, it seems that the closer an economic center is located to an affected district the greater is the reduction of FDI inflows observed, as discussed in Appendix A.4.
Second, for the sectoral analysis it is tested if the results change if one uses as shock weight sectoral employment share instead of sectoral GDP share. While the effects for manufacturing are in part comparable to the results discussed above, especially for construction, utilities, and services, different effects are found. A likely explanation is that GDP and employment represent different channels, as GDP, for instance, already includes productivity per worker. These different channels are the most likely rea-

sons for the differences. Both regression results are similar to the extent that upstream shocks tend to have a negative impact on FDI, whereas downstream shocks generally have a positive impact on FDI.

Limitations

Although they are often used in post-disaster analysis, Multi-Regional Input-Output Table (MRIO) tables suffer extensively from the fact that they are static. The linkage values used here are from 2016. Yet, a disaster might change supply-and demand networks within a country, which is ignored in this setting due to data limitations. Furthermore, alternative models such as the adaptive regional input-output models (ARIO) (Hallegatte, 2008) suggest improvements to identify production bottlenecks and to add more flexibility within production systems. Other recent methods also consider, for instance, the interaction between households, firms resources, and other agents in post-disaster recovery (see e.g. the ECAM model by Lehman et al. (2022)). A good overview of the shortcomings of the IO is provided by Galbusera and Giannopoulos (2018). Due to data limitations, the analysis is unfortunately based solely on a single MRIO table. Future studies should re-evaluate the results in a more dynamic setting.

1.4.3 Additional Robustness Checks

In addition to the previously mentioned robustness checks, it is also tested if the re-aggregation of districts before the decentralization policy in Indonesia induces a bias (Appendix A.2). This does not seem to be the case. Furthermore, multiple checks are conducted on the regressors included in the main approach. Appendix A.7 shows how reliable the shock metric is by re-running the baseline regressions with one shock or the other. While the shocks individually seem to matter, together the effect becomes stronger, highlighting the importance of considering both labor and capital in disaster studies. In terms of controls, nighttime-light overall tends to sufficiently proxy GDP in the case of Indonesia (Appendix A.8) except for one outlier province. Furthermore, the classification of economic centers using nighttime light seems to accurately represent the location of cities in Indonesia. In this section, it is also tested how reliable the main disaster data of Desinventar is. Overall, the data is strongly comparable to alternative datasets.

A last remark refers to additional checks on other disaster types. Appendix A.5 and A.6 test the regressions for flood events and show general effects using other disaster categories such as storms or hurricanes. Most districts are treated more than once with flood events; thus, to obtain a sufficiently large treatment sample, one has to allow for treatment reversal. In the appendix, it is found that, similar to earthquakes, flood events tend to reduce FDI inflows temporarily in the treatment year without long term effects. Sector-wise, fewer effects are found than for earthquakes, likely due to the lower destruction potential and better disaster preparedness in comparison to earthquakes (Neise and Diez, 2019). In terms of disaster categories, geophysical and climatological events (forest fires and droughts) tend to have the most severe impacts on FDI inflows, whereas hydrological and meteorological events only play a minor role. Yet this analysis is distinctly limited, so the results disaster category-wise should be taken with caution in their extent.

1.5 Discussion

In this section, a brief discussion is presented putting the found effects into perspective. The article documents significant temporary reductions of FDI inflows in affected areas with rebounding to pre-treatment levels within a year. The effects are stronger if neighboring districts are affected and upstream industries in within-country supply chains are affected. It also seems that the anticipation of the shock might matter. Exposure to disasters might lead to infrastructure destruction and blackouts, which can prolong recovery periods (Hallegatte, 2019). Therefore, the one-year fluctuation might be surprising, given how slowly infrastructure can be replaced. However, the recovery time from the shock in terms of FDI inflow is not that uncommon in the literature. Boehm, Flaaen, and Pandalai-Nayar (2019) documents for the Tohoku earthquake in Japan that production returns to pre-event trends within a year. Similar rapid recovery is documented in Inoue and Todo (2019), also for firms in Japan. Brucal and Mathews (2021) summarize studies as recovery periods of one to seven years after shocks. FDI flows might respond more slowly than direct firm-level responses given that investment decisions are revised. There might be pull factors after disasters attracting foreign investments (e.g. disaster aid or reduced tax burden) or push factors (e.g. high resilience or corporate opportunism) that might drive the quick rebound. Briefly, four of these factors are discussed.

The Government's Response The response of the government might offset some of the negative effects of the disaster. This can be through for the labor market response, as shown for the example of Chile and earthquakes (Jiménez Martínez, Jiménez Martínez, and Romero-Jarén (2020)) or through support for firm recovery (De Mel, McKenzie, and Woodruff, 2012). In such cases, the government response might attract foreign investment. The expenditure of the Indonesian government for disaster recovery is non-negligible. Between 2014 and 2018 around 1.4 to 1.9 percent of all government expenditure in Indonesia was dedicated to the recovery of shocks (The World Bank, 2021a). Governmental responses may also benefit foreign companies, potentially leading to a faster recovery. The government aid response is unlikely to be the primary driving force behind the observed effects for two reasons. First, it does not appear that firm recovery is at the center of the response to disaster risk in Indonesia. A study carried out by the World Bank (The World Bank, 2021b, p.12) shows that the implementation of approved reconstruction projects often takes more than a year to be decided and another to be initiated. This delayed funding is likely known to foreign firms. Subnational expenditures related to disasters between 2014 and 2018 primarily focused on reducing the risk of a disaster before it occurred, rather than rebuilding infrastructure (The World Bank, 2021b, p.74). Thus, firm recovery does not appear to be the focus of the government's response. Second, despite substantial advances in the government's role, disaster risk management and the connected government response became only politically important after the Presidential Regulation of 2008. Then a dedicated disaster risk agency, BNPB, was established. Assuming that it took some time for the BNPB to become fully effective, only a few years of covered study periods saw situations with a government response. Robustness tests revealed that the effects are present throughout the panel, providing some descriptive evidence that the government response may not be the primary driver of the rebound.

Tax Incentives Another, yet connected, point are tax incentives. Foreign firms may be incentivized to locate in special economic zones (SEZ) and benefit from free-trade zones or other tax benefits. In post-disaster environments, public policy might allow tax breaks for foreign firms to stimulate local economic growth. To our knowledge, the post-disaster relief from tax incentives is an understudied literature. Pan, Huang, and Jin (2024) study the effects of natural disasters on the tax burden on firms on the case of Chinese energy firms. They find that the corporate taxes are partially relieved to support impacted firms. In 2020, there were only 12 special economic zones OECD (2020) in all 416 districts, so the overall contribution to all districts is relatively small.

However, SEZs and tax incentives might influence the allocation of foreign firms to areas with certain risk profiles. More studies are needed to determine whether and how this selection in future treatments might influence responses from foreign firms. Regarding tax reliefs after disasters, standard corporate taxation is mandated by the Directorate General of Taxes (DGT) at the country level. At the local or provincial level, there appear to be at least no corporate income taxes that can be reduced (PwC, 2024). However, it remains unknown to us whether other non-income tax benefits can be decided on a district level. If they are decided at the province level, they are absorbed by the province-year fixed effects. Thus, the fraction of special economic zones is minor to the number of districts in Indonesia, yet we cannot with certainty exclude the possibility that some fraction of the quick rebound effect is due to local non-income tax breaks.

Firms opportunism Foreign firms could also be better connected to local politicians, thus receiving preferential distribution of disaster relief aid. Corruption and rent-seeking by firms in Indonesia have been studied in the literature (see e.g. Henderson and Kuncoro, 2011), thus it is a concern. While this article attempts to study the extent to which clientelism plays a role, no reliable data on district-level political clientelism or precise disaster disbursement were found. Thus, we acknowledge that there remains the potential that part of the quick recovery might be due to the involvement of foreign firms in local politics.

The Foreign Firm's Network Another potential explanation for the relatively quick recovery could be the integration of foreign firms in networks. The study has focused extensively on what happens if shocks occur along the supply chain. But while the shock might harm a foreign firm through a supply chain, foreign firms have distinct advantages. First, large firms might be able to join national conglomerates. This might also be true for foreign companies that enter domestic markets and aim to integrate themselves faster into national supply chains. Cole et al. (2019) provide an example for large Japanese firms, where large firms that join steel conglomerates had faster access to recovery funds. Given the discussion above, it is likely that recovery funds may not be the primary driver of the result. However, having denser conglomerates might nonetheless be beneficial. Foreign firms might be able to quickly adjust production by substituting for alternative suppliers. Input substitution would thus increase resilience (Dormady, Roa-Henriquez, and Rose, 2019), potentially accelerating recovery. More firm-specific input-output data would be needed to test this theory. Second, foreign firms might have the advantage of having quicker sales of their products to

foreign customers should local customers be affected. This is documented for Chinese firms that were able to target foreign markets after typhoons (Elliott et al., 2019). Both mechanisms, the stronger integration with large domestic firms and the ability to easily access foreign markets, could lead to a quick response of the FDI inflows.

Finally, investors might also refrain from direct investing and wait for the government's response before mobilizing money. Rapid recovery, if present, could serve as a signal to investors about resilience, leading to delayed but significant reinvestment in affected districts. The perception of risk and investors' sentiment after earthquakes have been found to be a contributing factor in investment decisions after a shock (see e.g. Shan and Gong, 2012). It remains not clear what is the main driving force behind the quick rebound. While the recovery time is not unrealistic, the conjectured hypothesis focuses rather on the signaling of recovery and the potential large resilience of foreign firms to absorb the shock (so rather push factors). Firms may delay investments and reinvest once the uncertainty is partially resolved. This would also be more in line with the wait-and-see concept in the literature on investment under uncertainty (Stokey, 2016). The government response and the tax incentives provided could explain small parts of the response, but are suspected not to be the main drivers. More research is needed to confirm this conclusion.

1.6 Conclusion

Previous literature dealing with the impact of disasters on foreign investment is mostly based on cross-country studies or firm-level data. Only a few papers focus on the impact of FDI within a country. We contribute to the literature by answering questions on the temporal, the sectoral and to a limited extent on the spatial dynamics of FDI inflows following the occurrence of earthquake exposure in Indonesia's economic centers between 2003 and 2019. For this reason, we use a shock variable accounting for labor and capital shocks to foreign firms. Using dynamic difference-in-difference models it is found that quite evidently in the case of Indonesian districts, FDI inflows tend to be reduced up to 90% in the year following a disaster event before returning to pre-disaster levels. The effect is mainly dominant in districts with a lower ex-ante risk as the disaster might be less anticipated, as well as in districts with a higher initial FDI stock, where each disaster can cause relatively more damage. These effects are robust against a set of alternative specifications, including recent methods taken from the difference-in-difference literature. In terms of spatial effects, it is found that

effects on FDI inflow are aggravated once a bordering district is also affected. A potential reason might be the supplier-customer network and the integration of foreign firms within that network. Using a multi-regional input-output table, we find that the strongest effects are found within manufacturing, which is also the sector that receives the most FDI in Indonesia. Hereby, especially upstream shocks, so shocks in the supply side towards foreign firms, tend to reduce FDI inflow substantially. Downstream shocks, in contrast, might benefit foreign firms, as one can see a rise in FDI inflow in most sectors in the year after the disaster. Construction seems to be one of the sectors benefiting from earthquake events, especially if they take place within the same sector or downstream in the supply chain. Thus, overall, one might carefully claim that the negative effects of FDI inflows are caused by intermediate disruptions in the supply of goods, which are necessary for production and predominantly play a role within regional production networks. This story would be in line with the prediction of Kato and Okubo (2022), who state in their paper that foreign firms are likely to remain in domestic markets if these firms are tightly linked with local suppliers.

The overall finding of the reduction in FDI inflow is consistent with most papers in the literature such as Friedt and Toner-Rodgers (2022) or Batala et al. (2021). Yet instead of a long term reduction, similar to the paper by Doytch (2019), we observe a rebound effect for FDI in the years following a disaster. This difference in long-term effects in the literature also contours one of the limitations of this paper: its external validity. Countries are globally unequally affected by disasters, specifically now through climate risk. On a provincial or district scale, there is a similar heterogeneity of disaster exposure. Each disaster might change a country's, province's or district's growth trajectory, which makes it econometrically complicated to compare multiple districts with each other. So, externally speaking, such assessments on the effects of disasters should be repeated on a country-by-country basis.

The results have real-life implications. Although FDI inflow declines in the aftermath of earthquakes, they do not appear to substantially reallocate or withdraw from the Indonesian markets in the years following. This resilience may be due in part to the ability of the firms to adapt and integrate into domestic supply chains. The heterogeneity analysis revealed that anticipation is a key factor in determining the response to the shocks. Policy makers should therefore focus on transparent communication of location-specific hazard risks. When, for example, the awareness of tsunami risk between Chile, Japan and Indonesia is compared, the local population appears to be less aware of tsunami risks in Indonesia than in the other two countries (Esteban et al.,

2013). This could also be the case for foreign investors who have never personally experienced earthquake-related risks. Thus, education on geophysical and increasingly climate risks is important. Furthermore, a clear and integrated system for the government's response to these shocks, implemented in a timely and efficient manner, could further mitigate fluctuations in FDI inflows. Although not the focus of this study, targeted government incentives post-disaster could have the potential to further mitigate negative effects. Further research is needed to support this claim.

Future research should extend the scope of this study to other disaster categories while distinguishing issues such as the expectation of exposure to a disaster as well as the supply-chain linkages of individual sectors. Thereby, the research shall include more dynamically weighted linkage models, that take bottleneck productions into account, as for instance described in Hallegatte (2008). A theoretical model that takes investors investment decision under uncertainty combined with indirect sector-linkages as in Kato and Okubo (2022) would also be a promising advancement. Lastly, future research should take a perspective of multi-hazard disaster risk as climate change makes connected disasters more likely (e.g. floods and extreme rainfall induced landslides). Papers with a focus on a single disaster type tend to fall short in accounting for the complexity of the repercussions of disasters such as compound events. However, on a local level, differences between different disaster categories matter, as disaster risk might be unequally distributed within sub-national districts. We only contoured this by including, besides earthquakes, earthquake-induced tsunamis (yet excluding earthquake-induced landslides).

In recent years, the literature on disaster risk has seen a renaissance. Disaster risk, especially through the lens of climate-related disaster risk, plays a key role in adapting countries to the long-lasting effects of anthropogenic climate change. Given the likelihood of increasing disaster frequency in the upcoming decades, it is important to prepare stakeholders, including local communities, and foreign firms, for the changes and challenges they might increasingly face. Research on historic data can serve as a helpful tool here. Findings resulting from such research can help the public and private sector to develop the best strategies for adaptation against disasters in order to minimize the loss of capital, safeguard the development progress of the last decades, and last but not least to minimize the impact on the livelihood of millions of people.

A Appendix to Chapter 1

A.1 Figures & Tables

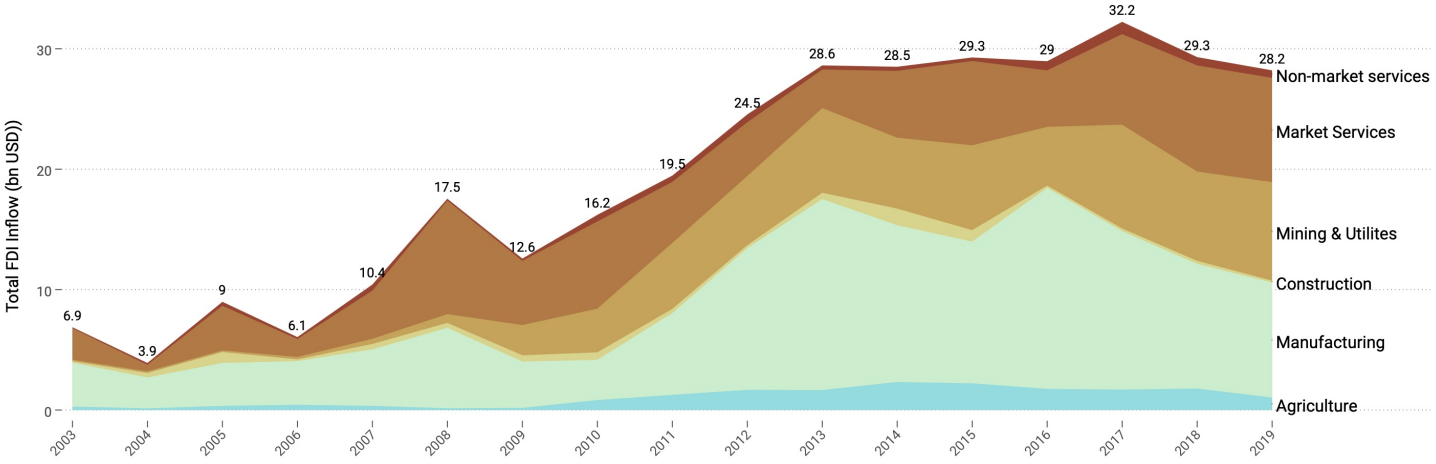

Fig. A.1. Development of FDI Inflows Into Indonesia (By year and sector)

Note: The data is based on Indonesian Investment Co-ordinating Board (BKPM) (2021). Author's own assignement.

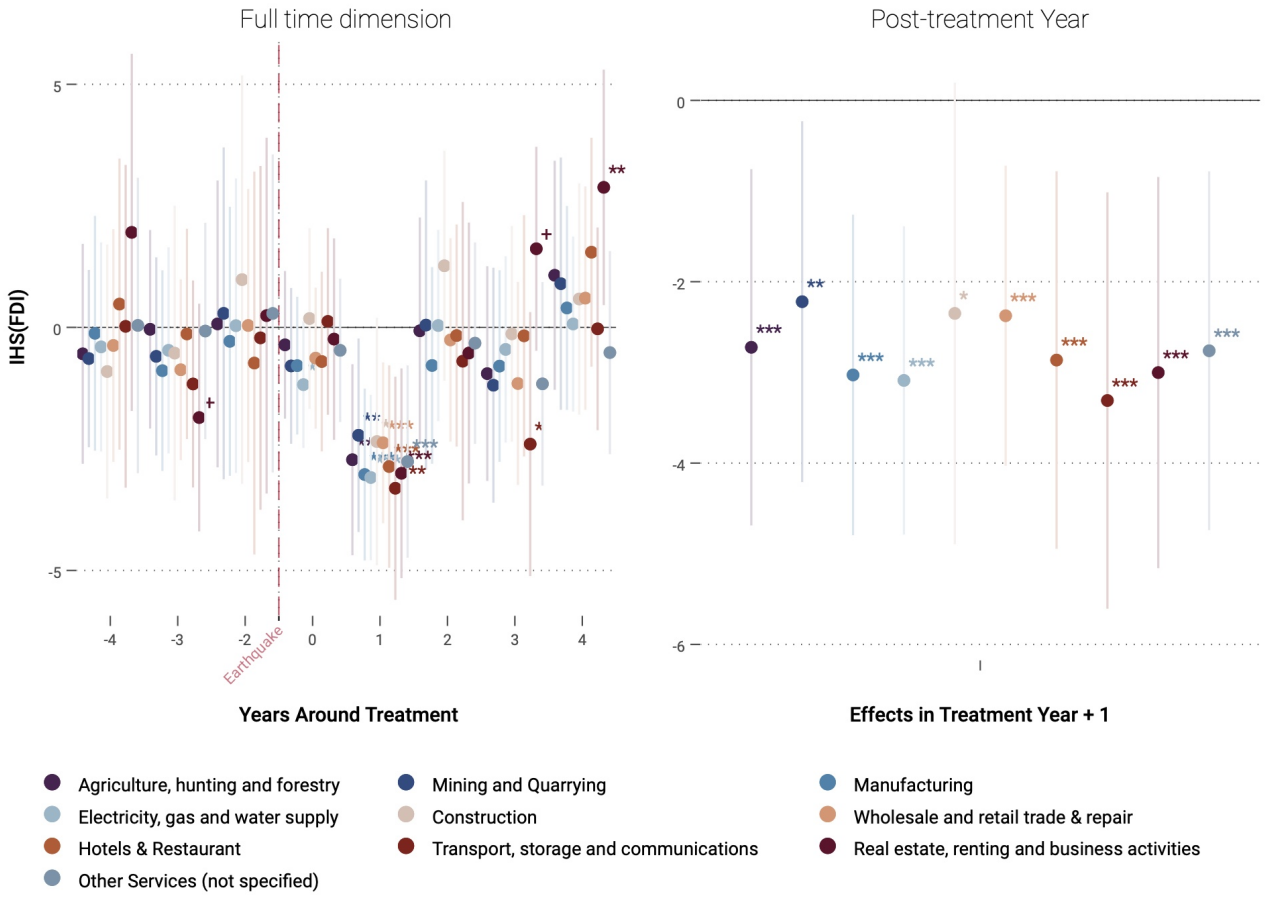

Fig. A.2. Effects on FDI (By Sector)l

Note: The two panels show the effects of Equation 1.1 if one extends the model towards the sectoral annual FDI space. Classification of FDI sectors follows broadly ISIC Rev. 3 level 2. The bars depict the 95% intervals and stars indicate *** 1%, ** 5% and * 10% significance. The left panel shows the effect over time, whereas the right one shows the effect solely in the year after the earthquake event (yet using the same regression as before). The treated group includes 38 treated districts and excludes as before treatment reversal and simultaneously occurring disaster events.

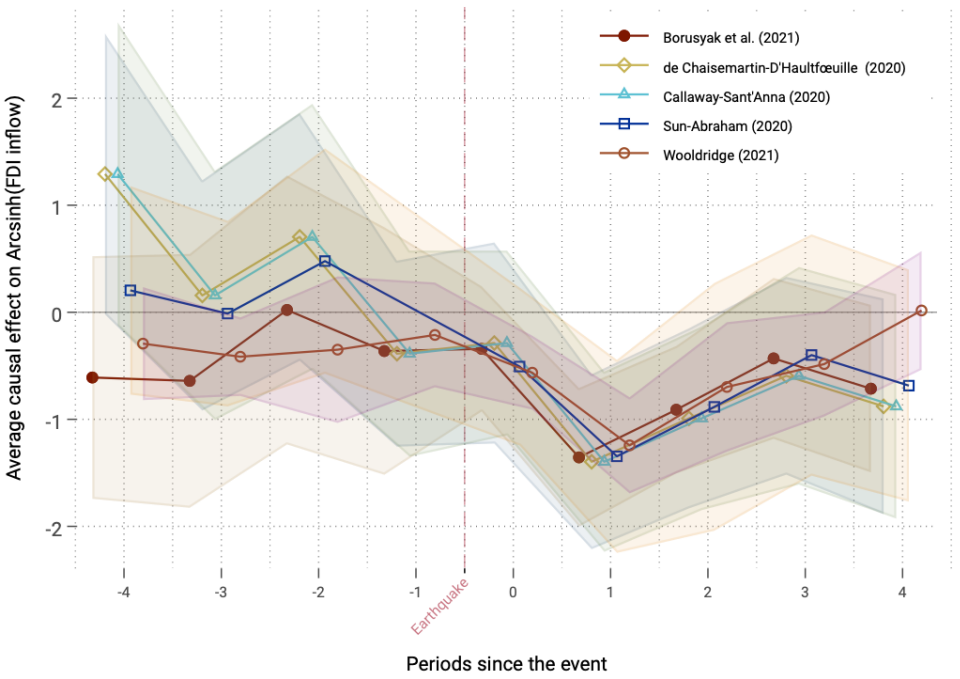

Fig. A.3. Comparison of Different Recent DiD estimators

Note: All included estimates allow staggered treatment design, yet exclude the possibility of treatment reversal. Here, each individual regression excludes multiple treatments within unit and uses only a sample of districts, which did not change names over the period (omitting 5% of the sample), which does not affect estimates substantially. The code is inspired by Borusyak (2022), , *five_estimators_example.do* (*Visited 14/02/2023*) and presents estimators by Sun and Abraham (2021), Sant'Anna and Zhao (2020), De Chaisemartin and d'Haultfoeuille (2020), Borusyak, Jaravel, and Spiess (2024), and Wooldridge (2021). The chosen CI bandwidth is the 10% threshold.

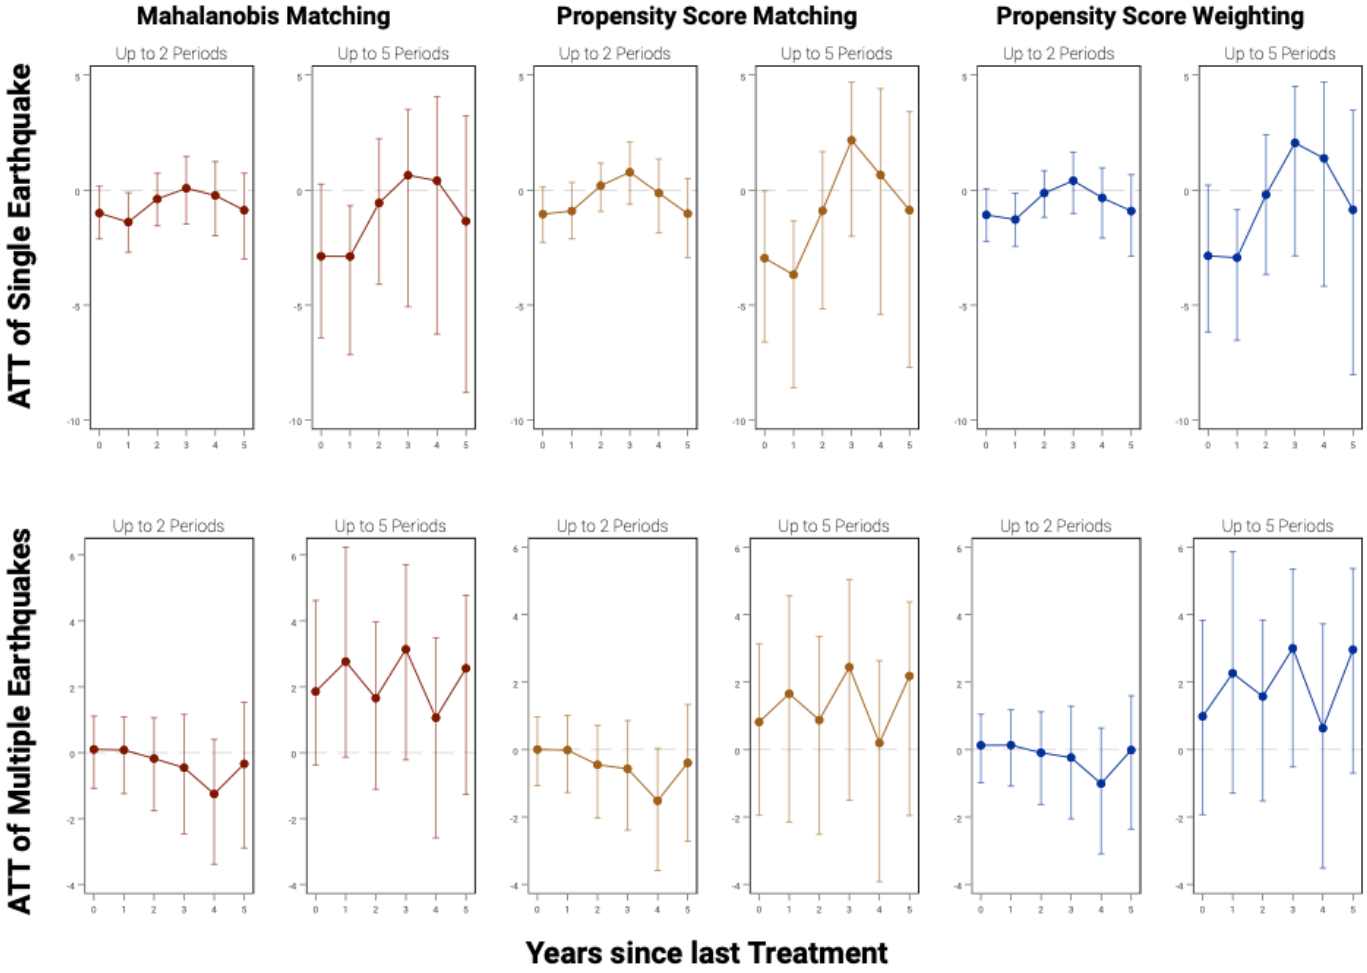

Fig. A.4. Dynamic Treatment Effects With Panel Matching

Note: The panels show the effect of earthquake exposure onto FDI inflows, where control units show a similar treatment pattern and are matched. This method is based on Imai, Kim, and Wang (2023). The upper panel excludes districts with more than a single earthquake recorded, and the lower panel excludes districts with only a single earthquake recorded. The three columns show three different matching algorithms for building a control group. The different columns show how many pre-treatment periods are used to build a matching control group. Matching covariates are based on the last three lags of nighttime-light, three lags of arcsin FDI inflow, the current log of population and the time-invariant composite hazard risk index. All matched control groups have a sample size of 10 and standard errors are 1000 times bootstrapped. The confidence intervals shown correspond to a 10% threshold.

Table A.1
Overview on Variables

Variable Name	Source	Modification	Description/Reasoning
Main Controls			
Harmonized Nighttime Light	Chen, Yu, et al., 2021	Averaged on district-level and Arcsin-transformed	Proxy for economic activity, especially in rural areas (see e.g. Sutton, Elvidge, and Ghosh, 2007; Hu and Yao, 2022)
Total Population	The World Bank (2022b)	Logarithmic-transformed	Might attract foreign investment through labor market potential and higher infrastructure
Extended Controls			
Annual Manufacturing share (% on district GDP)	The World Bank (2022b)	Based on sectoral GDP in million IDR (domestic currency). Is only available from 2010 onwards. Before 2010 value imputed based on 2010 values.	Incentivizes FDI through tightly-linked local supply-chains
Domestic Investment Flows	Indonesian Investment Co-ordinating Board (BKPM) (2021)	Arcsin-transformed	Might facilitate foreign greenfield investment, if they create local supply chain networks towards foreign company
World Bank Foreign Aid	AidData GeoQuery (Goodman et al., 2019)	Arcsin-transformed	Might attract foreign investment e.g. through infrastructure improvements
Bartik-like Migration Shift-Share	Share:Minnesota Population Center and Statistics Indonesia, 2020 and Shift:UN Economic and Social Affairs (Population Division) (2019)	Interaction of: People being foreign born in each district with five-year migrant stocks. Standardized by z-score	Following Javorcik et al. (2011) and Kugler and Rapoport (2007), migration might have an impact on FDI
Poverty Rate (% of Population)	The World Bank (2022b)	-	Alternative proxy for economic situation of a district
Morbidity Rate (in %)	The World Bank (2022b)	-	Alternative proxy for economic situation of a district
Population in Labor Force	The World Bank (2022b)	Logarithmic transformed	Control for the labor market potential
Other data			
Multi-regional I/O Table (2016)	Badan Pusat Statistik, 2021	-	To compute sectoral up-/downward linkages within and across provinces.

Table A.2
Event-Study Type Difference-in-Differences

	(1) Asinh(FDI)	(2) Asinh(FDI)	(3) Asinh(FDI)
Treatment Year	-0.725	-0.843	-0.548
	[-2.035,0.586]	[-2.260,0.573]	[-1.975,0.880]
Treatment Year +1	-2.734***	-2.854***	-2.741***
	[-4.387,-1.081]	[-4.521,-1.187]	[-4.312,-1.170]
Treatment Year +2	-0.491	-0.524	-0.121
	[-2.399,1.416]	[-2.441,1.394]	[-2.103,1.862]
Treatment Year +3	-0.751	-0.771	-0.428
	[-2.533,1.031]	[-2.553,1.012]	[-2.107,1.251]
Treatment Year +4	0.520	0.536	0.683
	[-1.430,2.469]	[-1.394,2.466]	[-1.325,2.692]
Main controls			
Log(mean homogenized NTL)(Lag +1)	0.297**	0.301***	0.229*
	[0.0710,0.522]	[0.0756,0.527]	[-0.00255,0.461]
Log(population (imputed after 2014)) (Lag +1)	0.342	0.403	-0.0898
	[-0.754,1.437]	[-0.724,1.531]	[-1.757,1.577]
Extended controls			
IHS(Domestic Investment in '000USD)			0.0848***
			[0.0565,0.113]
Poverty Rate (in % of population) (Lag +1)			-0.0795**
			[-0.154,-0.00471]
Morbidity Rate (in %) (Lag +1)			0.00958
			[-0.0173,0.0364]
Bartik-like Migration Control			687.0
			[-362.2,1736.2]
Manufacturing Share on Total GDP (in %)			16.40**
			[2.022,30.78]
Observations	4838	4794	4692
R^2	0.676	0.676	0.681
Districts in treatment group	48	47	47
Year FE	✓	✓	✓
District FE	✓	✓	✓
Province × Year FE	✓	✓	✓

Note: *<0.1, **<0.05 and ***<0.01 p-value. All regressions also contain pre-trends, which are excluded in the output. Bartik-like control is the initial share of migrants in 2000 interacted with 5-year national migration outflow. All models exclude multiple treatments and simultaneously occurring non-earthquake-related disasters. Model (2) and (3) additionally exclude districts with earthquakes between 1997-2002. All errors are clustered at the district level.

Table A.3
Dynamic Difference-in-Differences: Space

	(1) Shock: Labor+Capital	(2) Shock: Labor	(3) Shock: Capital
Base: No shock X No Neighbor Shock			
No Shock x Neighbor Shock	0.616	-0.368	0.702
	[-0.411,1.642]	[-0.954,0.217]	[-0.382,1.786]
Own Shock x No Neighbor Shock	0.128	-0.749	-0.470
	[-2.008,2.263]	[-1.884,0.387]	[-3.043,2.103]
Own Shock x Neighbor Shock	-0.160	0.0950	0.635
	[-1.630,1.310]	[-0.910,1.100]	[-0.778,2.049]
Base: No shock X No Neighbor Shock (Lag +1)			
No Shock x Neighbor Shock	-0.446	-0.739**	-0.361
	[-1.628,0.735]	[-1.331,-0.146]	[-1.342,0.619]
Own Shock x No Neighbor Shock	-2.932***	-1.213**	0.0337
	[-4.995,-0.869]	[-2.345,-0.0796]	[-1.608,1.675]
Own Shock x Neighbor Shock	-2.834***	-1.324*	-1.036
	[-4.740,-0.928]	[-2.784,0.137]	[-2.620,0.548]
Base: No shock X No Neighbor Shock (Lag +2)			
No Shock x Neighbor Shock	-0.822	-0.768**	-0.189
	[-2.104,0.459]	[-1.424,-0.113]	[-1.291,0.912]
Own Shock x No Neighbor Shock	-1.606	-1.494**	-0.144
	[-5.133,1.920]	[-2.749,-0.238]	[-1.468,1.180]
Own Shock x Neighbor Shock	-0.840	-1.138	-1.239
	[-2.988,1.308]	[-2.943,0.667]	[-2.849,0.372]
N	4071	4071	4071
R^2	0.674	0.674	0.673
Baseline Controls	Yes	Yes	Yes
Time FE	✓	✓	✓
District FE	✓	✓	✓
Province× Year FE	✓	✓	✓

Note: *<0.1, **<0.05 and ***<0.01 p-value. Dependent variable: IHS(FDI inflows). Each model includes log NTL and log population as Lag +1 term. Both models exclude treatment reversal and simultaneous unfolding other disasters. Labor Shock is defined if any person is recorded to be affected by disaster. Capital Shock is defined as peak ground motion above level V. The sample chosen excludes districts, which divided themselves after 2003 as here the centroids will be wrongly allocated.

Table A.4
Descriptive Statistics on Shock variables

	Sum	Mean	SD	Min	Max	N
Sectoral shock as % of GDP	9,067	1.58	9.66	0	97	5,746
Total Shock (Upstream)	4,326	0.75	4.22	0	87	5,746
Total Shock (Downstream)	7,132	1.24	5.66	0	102	5,746
Shock (Downstream) in Within Province	3,617	0.63	4.69	0	102	5,746
Shock (Downstream) in Rest of Indonesia	3,515	0.61	3.13	0	71	5,746
Shock (Upstream) Within Province	3,495	0.61	4.10	0	87	5,746
Shock (Upstream) in Rest of Indonesia	831	0.14	0.50	0	24	5,746

Note: Shocks are computed as interaction of sectoral GDP of districts within a province affected by any earthquake together with associated linkage value from the Leontief inverse matrix.

A.2 Indonesia's Decentralization Policy

Indonesia has decentralized many districts following the introduction of a decentralization policy in 2001. Thereby, the FDI inflow variable potentially accounts for different districts in different years, which could cause the earthquake shocks and FDI inflow data to be mismatched.

To avoid that problem, the paper standardizes the identification of districts based on their original 2003 names and identifiers. In other words, if a district is divided in the years following 2003, then it is re-aggregated to the name and the unique census identifier of 2003, the baseline year, to get a consistent level of observation over time. Out of the included 436 districts, this concerns approximately 51 districts, or 11% of the sample, that changed their names after 2003. The aggregation serves the purpose of not losing this data, yet it might be that thereby we overestimate the effects within these districts. Also, economic center identification might be inaccurate as it is done on the 2015 identifiers. By excluding the 51 districts from the analysis, one reduces the treatment sample by 31 of the 127 earthquakes recorded. This determines if and to what extent these districts, if excluded, bias the main results. Table A.5 shows that the results are pretty much comparable to the baseline results. Hereby, each uneven column includes still name switchers, whereas all even columns completely exclude districts that have switched their names. The coefficients of the controls are comparable. Excluding this subsample will matter in the case of extreme outlier events, precisely on occasions where there are ex-ante fewer observations in the treatment group, as each additional lost observation might substantially pull the regressions in a particular direction.

A.3 Heterogeneity Analysis of Empirical Strategy

Introduction of the Problem

The chosen empirical strategy is a dynamic two-way fixed-effects model that accounts for treatment heterogeneity by accounting for the time since the last exposure. Similar to classic difference-in-difference designs, this two-way fixed effects model relies on the parallel trend assumption. In research, this assumption is usually tested by considering pre-trends, yet given that, by nature, earthquake episodes are random in time, anticipation effects are unlikely in this study. However, one still finds that treated districts vary from control districts along covariates, which puts the parallel trend assumption under threat. To this end, in the main paper and in the following robustness

Table A.5
Impact on FDI Inflow Excluding Switching Districts

	(1)	(2)	(3)	(4)	(5)	(6)
Dynamic Treatment Trends						
Event	-0.725	-0.708	-0.843	-0.869	-0.548	-0.685
	[-2.035,0.586]	[-2.012,0.597]	[-2.260,0.573]	[-2.268,0.530]	[-1.975,0.880]	[-2.118,0.749]
Event +1	-2.734***	-2.868***	-2.854***	-3.032***	-2.741***	-3.002***
	[-4.387,-1.081]	[-4.538,-1.197]	[-4.521,-1.187]	[-4.705,-1.358]	[-4.312,-1.170]	[-4.605,-1.399]
Event +2	-0.491	-0.831	-0.524	-0.894	-0.121	-0.407
	[-2.399,1.416]	[-2.926,1.264]	[-2.441,1.394]	[-3.001,1.213]	[-2.103,1.862]	[-2.630,1.815]
Event +3	-0.751	-0.710	-0.771	-0.754	-0.428	-0.342
	[-2.533,1.031]	[-2.445,1.025]	[-2.553,1.012]	[-2.486,0.978]	[-2.107,1.251]	[-2.024,1.340]
Event +4	0.520	0.617	0.536	0.612	0.683	0.696
	[-1.430,2.469]	[-1.368,2.602]	[-1.394,2.466]	[-1.350,2.573]	[-1.325,2.692]	[-1.341,2.732]
Baseline Controls						
IHS(NTL) (Lag 1)	0.297**	0.298**	0.301***	0.314***	0.229*	0.237*
	[0.0710,0.522]	[0.0604,0.535]	[0.0756,0.527]	[0.0758,0.553]	[-0.00255,0.461]	[-0.00801,0.481]
Log(Pop Density) (Lag 1)	0.342	0.183	0.403	0.247	-0.0898	-0.0744
	[-0.754,1.437]	[-0.974,1.340]	[-0.724,1.531]	[-0.972,1.466]	[-1.757,1.577]	[-1.948,1.799]
N	4838	4345	4794	4318	4692	4223
R^2	0.676	0.676	0.676	0.676	0.681	0.680
Districts Treated	48	45	47	44	47	44
Excl. Switching Districts	No	Yes	No	Yes	No	Yes
Excl. pre-2003 Events	No	No	Yes	Yes	Yes	Yes
Extended Controls	No	No	No	No	Yes	Yes
Time FE	✓	✓	✓	✓	✓	✓
District FE	✓	✓	✓	✓	✓	✓
Province × Year FE	✓	✓	✓	✓	✓	✓

Note: *<0.1, **<0.05 and ***<0.01 p-value. Dependent variable: IHS(FDI inflows). All models exclude treatment reversal and the simultaneous appearance of other disasters. All regressions also contain pre-trends, which are excluded in the output. Extended controls included are lagged log of domestic investment, log of persons in labor force, manufacturing share, a Bartik-like control for migration, and lagged values of log morbidity and poverty. All models exclude multiple treatments and simultaneous occurrences of non-earthquake-related disasters. Model (2) - (6) additionally exclude districts with earthquakes between 1997-2002. All errors are clustered on district level.

tests are conducted to evaluate this problem in this study.

Furthermore, other problems can arise with the chosen method: first, treatment timing varies across district groups. If one would only consider those districts that are exposed once to a stronger earthquake, one would end up with what is often referred to as a "staggered Difference-in-Difference adoption design". This would correspond to the treatment pattern shown in Panel (a) in Figure A.5. Choosing this pattern excludes districts with more than one earthquake event (around 8% of the total sample), so important information could be lost. The baseline model follows this setup. Panel (b) shows the true treatment pattern, which is empirically more difficult to address due to the possibility of treatment reversal. To correct for this staggered adoption design, one needs to weight different treated groups differently. If one does not account for the weighting, a bias arises (discussed, for instance, by Xu, 2023; De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Borusyak, Jaravel, and Spiess, 2024).

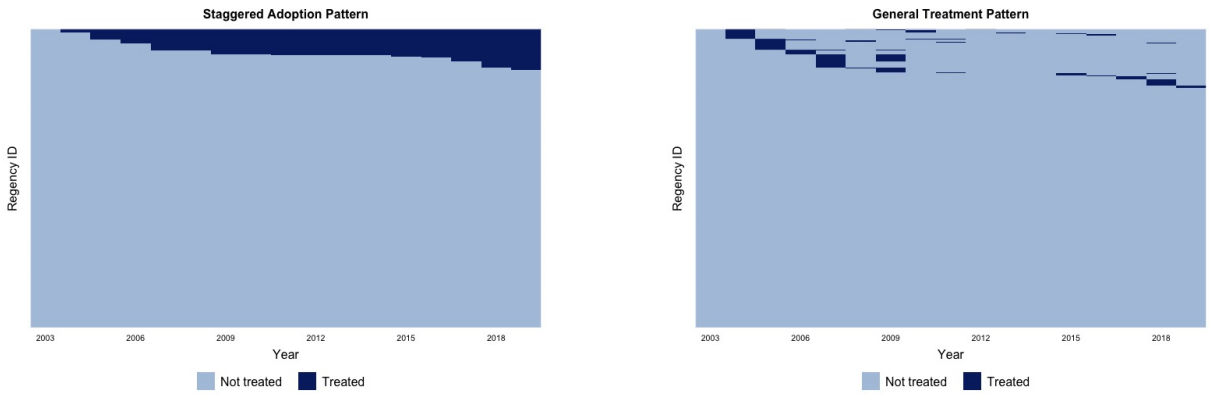

(a) Staggered Adoption

(b) General Pattern

Fig. A.5. Different DGPs

Note: The graphs have been generated using Xu and Licheng (2018) `panelView` command and the main data frame created for this study.

A second problem is that the treatment definition relies on the strict exogeneity assumption. Following Xu (2023) and Imai and Kim (2019) strict exogeneity relies on four main assumptions, namely the absence of confounding variables that vary over time, the absence of lagged dependent variables, the absence of an effect of past outcomes on the likelihood of being treated and lastly, that treatments from the past do not impact future outcomes. If these assumptions do not hold, then estimates are likely to be biased. To check for time-varying confounders, one usually includes dynamic treat-

ment effects, which we respect by differing between different timings since the last treatment. So far, it has not been considered how past outcomes impact current ones. Such a case could arise if foreign firms withdraw from local markets following a disaster. That past outcomes (FDI inflow) impact the treatment's existence is excluded by design in our case. To summarize, the carryover effects of past treatment on current outcomes and of past outcomes on current outcomes may potentially bias the estimates. Furthermore, in the chosen dynamic Two-Way Fixed Effects (TWFE) model it is likely that found effects are biased by effects of other periods (Sun and Abraham, 2021), which needs to be addressed.

Dynamic Heterogeneous Treatment Effects

This section provides a brief discussion of Figure A.3, mentioned in the main paper. Figure A.3[15], equivalent to the baseline model, shows the effects of a singular earthquake exposure excluding the possibility of treatment reversal A selection of recent difference-in-difference estimators from literature are used (Sun and Abraham, 2021; Sant'Anna and Zhao, 2020; De Chaisemartin and d'Haultfoeuille, 2020; Borusyak, Jaravel, and Spiess, 2024; Wooldridge, 2021). All models include only districts that have not been split over time in order not to overvalue FDI inflows in total. The mentioned papers address the negative weighting problem mentioned above in various ways. For instance, De Chaisemartin and d'Haultfoeuille (2020) suggest to use weighted average of two difference in difference estimators: the first compares districts switching from not being exposed to an earthquake to experiencing an earthquake with districts that have not experienced any in both periods. The second estimation compares districts that have experienced an earthquake at both points in time with those that switch between the two periods from not treated to treated.
All estimators used follow a similar pattern post-treatment. Following a strong earthquake exposure, one sees a persistent decrease in FDI inflows in the years following the disaster with the most consistent decline found in the DiD_l estimator by De Chaisemartin and d'Haultfoeuille (2020) and the Sant'Anna and Zhao (2020) estimator. In terms of significance, all methods find a statistically significant reduction in FDI inflows in periods 1 and 2 after disaster occurrence at a 10% level. The coefficient is similarly large as in the baseline regression. Overall, the effects found effects, confirm the baseline results.

[15] The code is inspired by Borusyak (2022) (*Visited 14/02/2023*)

Allowing For Treatment Reversal

As mentioned above, the baseline model excludes, by definition, districts that experience more than a single disaster shock during the study period. To investigate whether this exclusion substantially changes our results, the estimator suggested in De Chaisemartin and d'Haultfoeuille (2023) is applied. In this setup the treatment group consists of districts being treated once or multiple times as here treatment reversal is possible. The results, shown in Figure A.6 are broadly in line with the results of the baseline regression, whereby one observes a small "dip" in the year after the event, significant on a 10% level. The left figure shows the baseline setup including controls, yet not lagged controls, while the right model adds linear year and district as well as non-linear province-specific trends. In all cases, standard errors are 100 times bootstrapped, and especially in the right panel for some replications, the p-value is marginally above the 10% threshold. While there is an indication of a pre-trend occurring four years before the treatment, the treatment reversal approach is somewhat hampered by the limited time available for recovery before the recurrence of earthquakes, which could potentially introduce some contamination to the pre-treatment periods. The results broadly confirm that excluding the multiple treated units does not significantly alter the findings. However, multiple treated units might be fundamentally different from single treated ones, which is why the group below is considered separately.

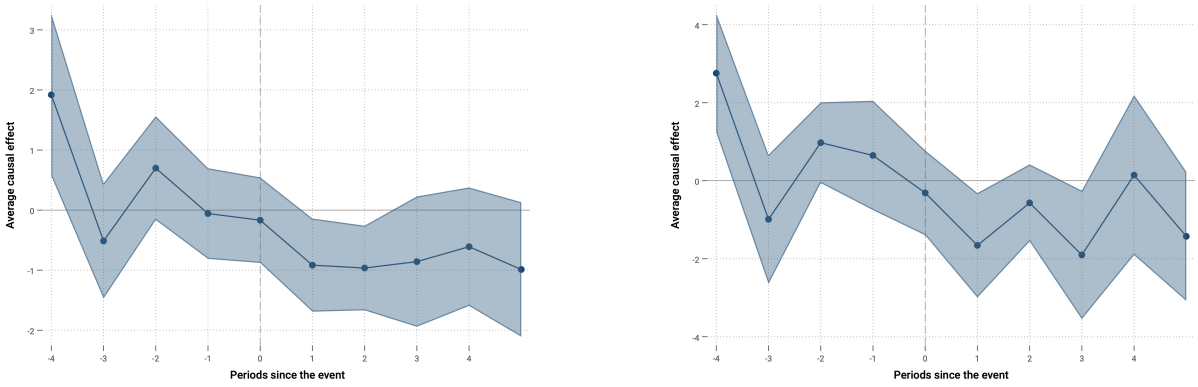

(a) Baseline Model **(b)** Adding different trends

Fig. A.6. Allowing for Treatment Reversal

Note: The graphs are based on De Chaisemartin and d'Haultfoeuille (2023) and plotted using Borusyak (2021). Panel (b) includes linear time and district trends, as well as non-linear province trends. Confidence intervals are based on a 10% level. Standard errors are clustered-bootstrapped 100 times.

Matching the control group

As discussed in the main paper and shown in Figure A.4 in the paper, the panel matching approach by Imai, Kim, and Wang (2023) is used to test if the baseline results hold after controlling for differences between treated and control units. The methodology essentially first finds for each treated units comparable units with a similar treatment pattern. For this purpose a window needs to be defined upon which the treatment pattern should match. To compare changes in the results, we set this window either for two years or five years. For robustness, we also tested the results using other windows. We set the size of the control group to 10 districts for each treated unit. Having identified the set of control districts, a matching (or weighting) procedure is used to minimize covariate balance between treated and control group. As set of covariates to be matched (or weighted) on we use three lags of nighttime light, as proxy of economic development, three lags of FDI inflow, as proxy of investment patterns, current logged population as well as composite hazard risk as mentioned in the heterogeneity part of the article. These variables should proxy broadly historic investor perception, measured by comparable investment flows, attractiveness in terms of labor force and economic strength as well as disaster exposure risk, measured by the composite disaster index. We compare three refinement methods: Mahalanobis matching, propensity score matching and propensity score weighting. When studying the balance of covariates before and after refinement of the control group, the weighting of the propensity score and the matching of the propensity score outperform the average matching of Mahalanobis. In terms of pre-treatment parallel trends across covariates, propensity score weighting outperforms both other methods for all covariates except nighttime light. However, night-light trends show some fluctuation in the covariance balance. Thus, while the risk of violation of the parallel trend assumption is substantially reduced with this approach, differences in economic growth trajectories cause some noise in the estimated effects. However, it cannot be excluded given that economic growth, measured by nighttime light depicts a major driver for foreign investment. After identifying a set of control units, refining them across covariates, a difference-in-difference method is applied, whereby in the shown results, standard errors are 1000 times bootstrapped. Looking at the upper panel of Figure A.4, the one where the treatment definition matches the one of Equation 1 of the paper, all setups show comparable results to the baseline model of our paper.
As mentioned before so far in this section, the results have focused on discussing the negative weighting issue as well as if the found effects change, if districts are included

that have been treated more than once. The matching approach for now has only been used to minimize the reliance on the parallel trend assumption, yet given that it matches on treatment history it can also be used to assess if the results of the baseline regression change if one only considers districts with multiple treatments recorded. This is done and shown in the lower panel of Figure A.4. In contrast to what has been tested before, the treatment group here consists only, not additionally, of districts that have seen more than a single disaster over the study period. In these districts, which make up around 7% of the total sample or around a third of the treatment sample, one might expect to see different responses. While one cannot say with certainty that these districts are the ones with the highest frequency[16] given their more frequent exposure to earthquake events it is likely that these are areas where investors and policymakers are more prone to price in earthquake risk. The method used in this setup corresponds to the description of the methodology above. No significant effects are observed. Using more years to select control units increases the magnitude of the coefficient, given that only very few treated units follow a similar treatment pattern over a five-year window. The results in combination with the baseline results might indicate that firms in districts with higher earthquake event frequency are more adapted or more resilient to the risk, causing them not to change their investment behavior after the occurrence of an event. Future studies will need to highlight more of the implications of such accumulating temporal exposure to earthquake shocks.

Assessing the Parallel Trend Assumption

A key concern of the methodology used it the potential violation of parallel trends, which is a problem commonly faced in event-studies. To evaluate at which point the significant effects found vanish if one deviates from the linear structure of parallel trends, we apply a methodology by Rambachan and Roth (2023). Thereby, we can assess the sensitivity of the violation of the necessary parallel trend assumption in this study. Figure A.7 plots the analysis based on the regression of Equation 1.1, where in red the original 95% confidence intervals of the first significant effect is shown followed by the tested robust confidence sets. The parameter M on the x-axis shows how the FDI inflow outcome variable is allowed to differ between districts with earthquake and without. If M equals zero, then the difference in trends of the counterfactuals has

[16] As earthquakes are random in time, it can be by chance that over the study period these districts have seen more than a single destructive earthquake event.

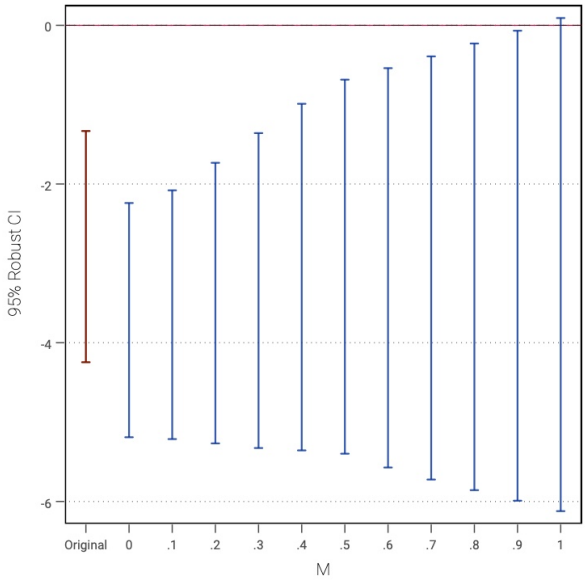

Fig. A.7. Credible Confidence Intervals: Testing for Non-Parallel Trend

Note: The shown graph is based on Rambachan and Roth (2023). Estimation is based on Equation 1.1. Post-treatment period is defined at the year after the treatment assuming that it takes time for the effect to appear. Sensitivity analysis based on smoothness restriction using second differences.

to be linear. If it is larger than 0, it can be nonlinear, so one allows that the pre-trend can change by no more than the value M between two successive years. The figure shows that the first significant effect found in the year after earthquake occurrence is robust against linear and non-linear violations of the parallel trends up to a value of 0.9. This provides evidence that the effect found in the year after treatment is not substantially driven by violations of parallel trends. For additional robustness, the exercise is repeated using the approach of Callaway and Sant'Anna (2021) as it accounts additionally for the negative weighting issue. The result is still robust against linear and nonlinear violations of the parallel trend, however with a breakdown value at already at 0.2.

Conclusion

The tests conducted here provide additional confidence that the negative weighting, as well as the selection of the treatment definition, do not fundamentally alter the baseline results of the article. Given the lower and differential trends of districts exposed to multiple earthquake events it seems reasonable to exclude them from the treatment definition. Matching on pre-treatment characteristics seems to not essentially change our results and testing for violation of the parallel trend assumption seems to indicate that the core results of the paper are robust against linear and standard non-linear

violations of the parallel trend assumption.

A.4 Reallocation Effects to Economic Centers

In the main paper, the impact of a district's FDI inflow on neighboring districts is tested. Yet, FDI might be predominantly allocated to economic centers.
In this section, not only is the economic center of each district used, but mainly the distance to the main economic center within each province, which is identified using nighttime light data. The main economic center within a province is, in essence, the district with the highest nighttime light intensity at the district-specific economic center. The classification of the main economic centers is validated with various statistical yearbooks. The distance from each economic center within a province to the closest affected district is then measured. Hereby, it is also allowed for this district to be located in a neighboring province, yet what matters is the distance to an affected area. Table A.6 shows the results. One can clearly see that proximity matters: the closer an economic center is located to an affected district, the larger the impact on FDI inflows, ranking from a 98% reduction within 50 km significant on a 1% level, down to 92% within 100 km, to 90% within 200 km. After that, no further significant effects were found. A key limitation of this model is the small number of treated economic centers, as there are only 32 provinces, and thereby only 32 within-province economic centers are available.

A.5 Effects By Floods

As shown in the paper's introduction, earthquakes and the associated tsunamis account for the vast majority of deaths and damages recorded in all disaster categories in Indonesia. In terms of population affected, however, floods seem to be the most relevant disaster category. Thus, in comparison to earthquakes, the same effects are analyzed similarly.

Identification

While formulas and models for geophysical ground motion are easily obtained for earthquakes, accounting for a physical measure of flood exposure is more complicated. We use extreme rainfall shocks as a proxy for flood risk. First, data from Abatzoglou

Table A.6
Impact on FDI Inflow in Economic Centers in Proximity to Disaster

	Distance to exposed district			
	(1) 50km	(2) 100km	(3) 200km	(4) 500km
Shock in distant district	-1.973*	-1.020	-1.198	-0.309
	[-3.965,0.0192]	[-3.065,1.024]	[-3.497,1.101]	[-6.118,5.500]
Shock in distant district (Lag +1)	-4.199***	-2.631**	-2.400*	-1.162
	[-6.650,-1.747]	[-4.816,-0.447]	[-5.047,0.247]	[-6.495,4.170]
Shock in distant district (Lag +2)	-0.257	-0.326	0.158	-0.223
	[-3.087,2.572]	[-2.751,2.100]	[-2.817,3.134]	[-10.52,10.08]
N	356	349	317	77
R^2	0.727	0.722	0.725	0.560
Economic Centers in Treated Group	7	9	7	1
Baseline Controls	Yes	Yes	Yes	Yes
Time FE	✓	✓	✓	✓
District FE	✓	✓	✓	✓
Province× Year FE	X	X	X	X

Note: *<0.1, **<0.05 and ***<0.01 p-value. Dependent variable: IHS(FDI inflows). Each model includes log NTL and log population as Lag +1 term. Both models exclude treatment reversal and the simultaneous occurrence of other disasters. Also, as centroids of each district are used, districts that switched names are systematically excluded.

et al. (2018) on monthly rainfall in mm since 1959 is used and overlaid with the 5x5km grid cells of Indonesia. As the time dimension in this paper is annual, the monthly data need to be re-aggregated on an annual level. For this purpose, the paper follows Felbermayr, Gröschl, et al. (2022) and the definition of a rolling window metric. This is calculated in the following way: First, the grid cell-month specific rainfall is standardized

$$\gamma^{\text{prec}}_{i,m,t} = \frac{x^{\text{prec}}_{i,m,t} - \bar{x}^{\text{prec}}_{i,m}}{\sigma^{\text{prec}}_{i,m}}$$

where i refers to grid cell identifier, m to month and t to year. For the calculation, the entire time period is used (since 1959), whereas the final dataset only includes the years 2003 to 2019. Next, only positive deviations are assumed, as only positive rainfall shocks are likely to be associated with flood events. To aggregate it on a yearly level, one cannot just average it over the year as this would omit dry and wet seasons within the year, whereby rainfall shocks, especially in wet seasons, are likely to cause more damage. So in the aggregation formula, each month is weighted with the remaining number of months within a given year, while the right-hand term is a weight for the

value in the next year:

$$D_{i,t}^{p} = \sum_{m=1}^{12} \left[\gamma_{i,m,t}^{p} \cdot \frac{12-m}{12} + \gamma_{i,m,t-1}^{p} \cdot \frac{m}{12} \right]$$

The value of D in grid cell i year t is then the weighted average of positive deviations from the long-run monthly mean. Again, nighttime light is used to identify economic centers. Then the rainfall shock metric is assigned to the economic center grid cells, if the rainfall metric is positive. Furthermore, using the historic risk maps from UNDRR (2017a) only economic centers are assumed to be treated with a continuous rainfall shock if the grid cells of the economic center are at historic risk of a 200-year flood. Then the mean value of these intensities in economic centers within each grid cell is finally used. So the rainfall shock measures the mean standardized rainfall shock over all historic monthly values in economic centers at risk within a district. The shock on the labor market follows the share of flood victims over the lagged district population. As before, both shocks are interacted and ranked.

There are several criticisms that can be raised regarding this methodology. Rainfall is an insufficient predictor of flood exposure, especially of riverine floods, as neighboring districts might be more affected by large rainfall upstream of the river. Also, the formula does not account for anthropogenic changes of the urban landscape over time, which might cause areas at risk to change. Nonetheless, for a robustness check, this shall suffice as a method.

Effect on Temporal Level

Figure A.8 shows the temporal effects of flood exposure. Panel (a) differs between the impact of using a top 10 ranked flood event and those within the top 100 ranking, where the ranking is based on the product of the labor market and the rainfall shock. The problem is that for the top 100 ranking only 38 events are included, as the entire rest has more than one treatment. Panel (a) shows that there appear to be no significant effects on FDI inflows through floods in the case of Indonesia. Yet, allowing for treatment reversal in Panel (b) using the DiD_l estimator by De Chaisemartin and d'Haultfoeuille (2023) one can see that FDI appears to be substantially reduced within the year of flood occurrence but not afterwards, so the effects appear earlier than in the case of earthquakes. Overall, the effects appear to be negative but not as obvious as for earthquakes since potentially flooding has other mechanisms through which it affects

foreign investment. An effect on a sectoral level is discussed next.

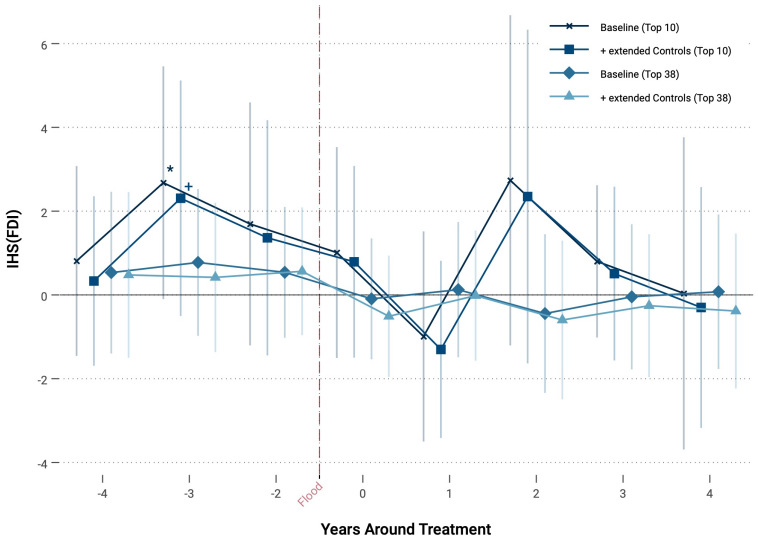

(a) Baseline Model

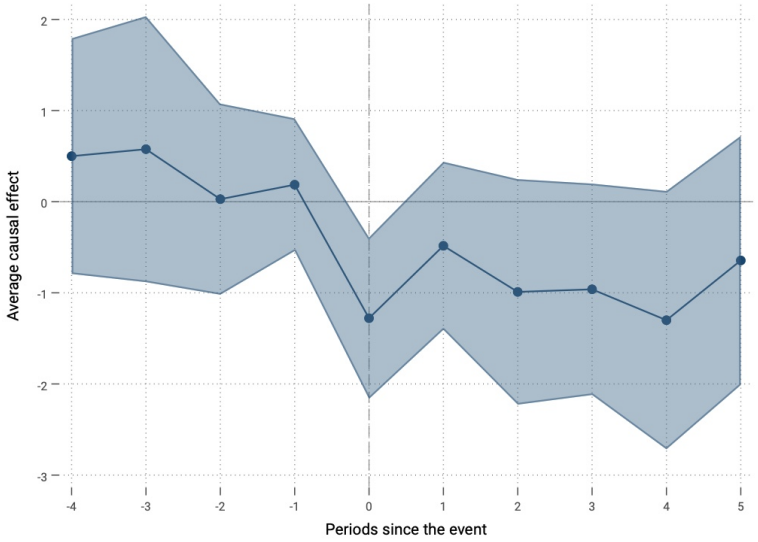

(b) Allowing for treatment reversal

Fig. A.8. Temporal Effects of Flood

Note: Graph (a) uses the main specification and distinguishes between Top 10 district-year events and Top 100 (out of which only 38 are included (as those districts with multiple treatments are in this specification excluded). The used estimation model is the one from Equation 1, so using the same core controls, including district, year and province-year FE and applies district-clustered standard errors. The graph (b) is based on De Chaisemartin and d'Haultfoeuille (2023) and plotted using Borusyak (2021). Panel (b) includes linear time and district trends and allows for treatment reversal. Confidence intervals are based on a 10% level.

Effect on Sectoral Level

Re-running the regression from equation 1 on a sectoral level for flood events, one finds mild positive effects of FDI in mining and stronger effects for wholesale and services. The differences in sectoral effects between flood and earthquake can be attempted to be explained through the distribution of the respective disaster hazards across Indonesia. Not all regions are affected by floods and earthquakes at once, which causes the sectoral FDI results to change.

Figure A.9 adds a network perspective by showing shocks through linkage effects once for extreme events (Panel (a)) and once for weaker events (Panel (b)). At first sight, one can see weaker effects through linkages in panel (b). Panel (a) shows that, surprisingly, through floods in most sectors, downward linkage shocks cause FDI inflows to reduce in the year after the disaster. As an exception, utilities show a significant rise in FDI inflows with a coefficient of 0.68. Upward linked shocks cause a reduction of FDI inflows only in the construction and mining sectors, yet hotels see a substantial rise in FDI inflows. Within-sector strong negative effects are found for utilities and small positive effects for mining and services.
In Panel (b) weaker flood within-sector-shocks cause FDI inflows into transport to decline and to mildly increase in real estate. Moving on to downstream shocks, real estate and agriculture are the only ones seeing a reduction in FDI inflows, yet agriculture sees an FDI increase through upstream shocks. Here again, it becomes clear that weak floods seem to take place in other regions than weaker earthquakes, which explains the differences in effects on a sectoral level.

Explanation of Differences

A key question that arises from this section is why the effects are so different, especially at the sector level, between earthquake and flood exposure. Temporally speaking, in both cases, after exposure to disaster, FDI inflows seem to reduce either immediately or with a delay of a few years. Yet, the origin of the effects might be different.
First, it is shown in the introduction that floods seem to have different effects on firms and populations than earthquakes. Floods tend to affect more people in their daily lives, whereas earthquakes are more destructive and deadly. For firms, domestic or foreign, they might adapt differently to either shock. Neise and Diez (2019), for instance, discuss how small and medium firms tend to adapt to flood risk. While you can adapt to flood risk for instance by investing in backup technologies, earthquakes and

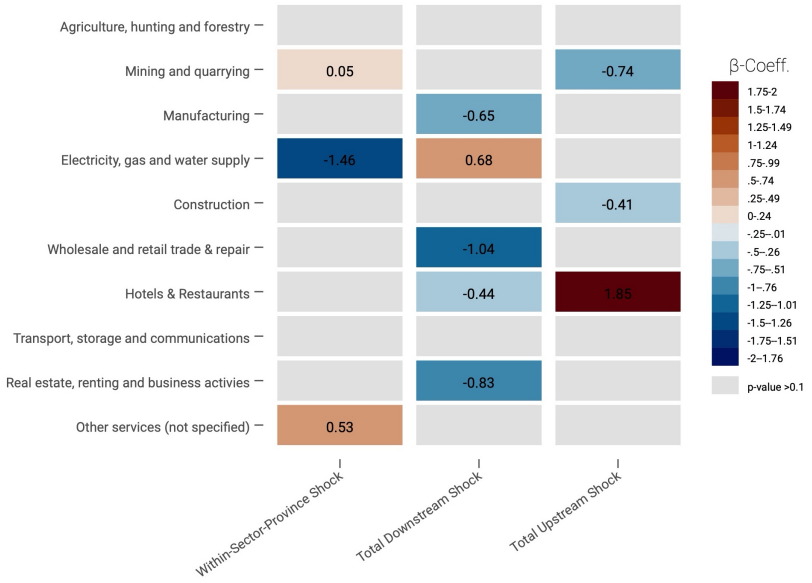

(a) Top 10 Events

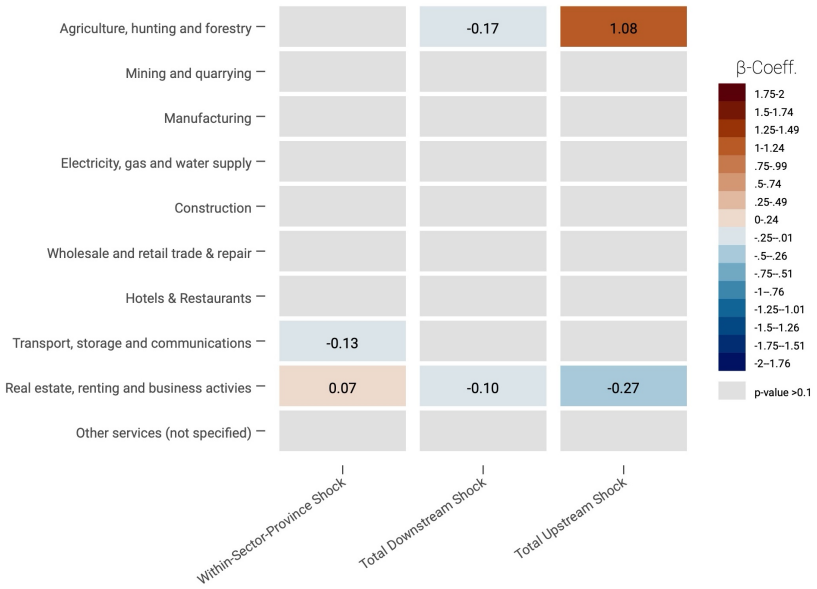

(b) Top 100 Events

Fig. A.9. Linkage Effects of Flood

Note: Values for β coefficients are shown. The two panels show the sectoral response of IHS-transformed FDI inflows based on a 1% flood-induced GDP shock. The first columns in both panels refer to a 1% shock occurring within the sector of FDI considered. The other two columns to up- and downstream shocks, described as Leontief-inverse matrix weighted GDP shocks. In contrast to the first panel, the second differs also by origin of the shock. Grey fields are below the 10% significance threshold. Estimations are based on Equation 3. The classification of top-events is based on the continuous function of the total shock and each regression includes lagged and logged nighttime-light and population values as core controls.

their associated tsunamis tend to cause more long-lasting damage to local economies, which would explain why the effects found vary.
Secondly, the expectation of disaster risk might vary between the two categories. Earthquakes are prone to occur in proximity to tectonic plates, while their associates, tsunamis, are more likely to affect low-land coastal regions. Floods also tend to originate in lowlands, and especially in proximity to rivers. Yet, due to anthropogenic climate change, floods induced by extreme rainfall events tend to become more frequent in the upcoming years, which is potentially not yet as deeply incorporated in the risk mitigation strategies of foreign firms. Lastly, as mentioned briefly before, the sectoral composi-

(a) Earthquake Risk and FDI

(b) Flood Risk and FDI

Fig. A.10. Spatial Distribution of Risk by Disaster Type

Note: Graphs are based on the OCHA's Regional Office for Asia and the Pacific (2021), data from Indonesian Investment Co-ordinating Board (BKPM) (2021) and hazard maps on 250-year return periods for earthquake and 200-year return periods for flood risk (UNDRR, 2017a; UNDRR, 2017b). Composition based on percentiles. In gray there are the 14 districts non-matched. The graph likely overestimates FDI slightly in 51 districts. Maps created using *bimap* command by A.Naqvi (2022).

tion of affected districts might vary between the two disaster categories. Figure A.10 plots in a bivariate map the FDI inflow against the calculated district-specific compos-

ite hazard index[17]. One can clearly see that areas (in blue) of high risk for earthquakes are located in different provinces than those for flood risk, which likely causes the sectoral effects to vary.
These are just a few of many reasons why floods might have different effects on foreign firms than earthquakes, yet nonetheless, it shows that flood events also do not generally contribute positively to FDI inflow, similar to earthquake events.

A.6 Impact of Other Disaster Types

The paper primarily focuses on the impact of earthquakes on FDI, as they pose a significant threat to human and economic life in Indonesia. In the previous section, it is more in-depth investigated what the effects of floods on FDI inflow by using another physical intensity proxy. This section only outlines the effects on FDI inflow, where disaster types are categorized into four major categories: geophysical disasters (earthquakes, tsunamis and volcanic activity), hydrological disasters (floods, landslides and tidal waves[18]), meteorological (storm and cyclone exposure) and lastly climatological disasters (drought and forest fires). All but geophysical events are described in the literature as being directly related to climate change. However, there is even a nascent literature on the impact of, for example, changes in tsunami propagation due to sea level rise (e.g. Sepúlveda et al., 2021) or how related meteorological events might cause additional stress on tectonic crusts (Steer et al., 2020). To analyze the differences between various disaster shocks on FDI inflows, quarterly FDI data are used to get as close as possible to the point in time when the disaster occurs. Here, this is possible, as for this study, there are no physical measures for the other disaster categories. Certainly, it would be possible to account for wind-speed models, volcanic ash fallout or wildfire spread using remote sensing and interpolation mechanisms, but this would exceed the scope of this paper. Thus, to measure disaster events, the Desinventar by the UNDRR (2022b) database is used. A disaster is considered a disaster if at least one person has died or been directly impacted by it. Admittedly, this is a very strong assumption, but in this way, the maximum number of disasters can be retained as a dynamic difference-in-difference specification used in the main part of the paper, whereby treatment reversal is excluded by definition. Most disasters occur more than once over a

[17] The risk index, as described in the respective section, is created by dividing a standardized grid cell-specific mean of each district into quartiles over the entire country
[18] Waves caused for instance by tidal gauges

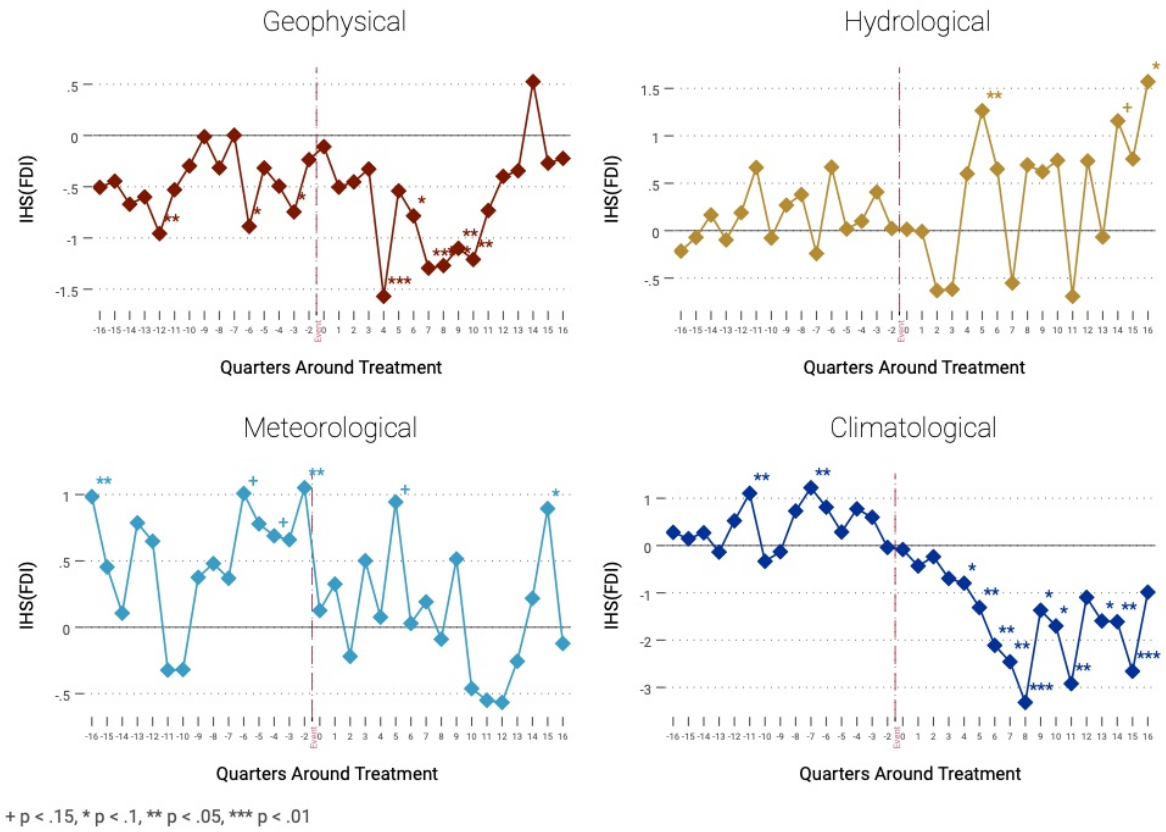

Fig. A.11. Dynamic Effects by Hazard Category

Note: Figure shows the effects of aggregated disaster events, classified by their characteristic, onto quarterly FDI inflows aggregated over sectors. Estimations are based on Equation 1 and are conducted at the district level. Included are all events by disaster category, where only a single event within a district in this category occurred (districts considered treated: 72 for geophysical, 44 for hydrological, 59 for meteorological and 63 for climatological events). Other disaster events within a district in the same quarter are excluded. A district is considered treated, if at least a single person has died and/or is affected by a disaster.

17-year period, which is why a very low threshold of at least a single person affected is used to include a larger sample of treated units.

Figure A.11 shows the results of the regressions. One can see that, similar to the models used in the main part, even by including other volcanic activity and other tsunamis, geophysical events cause a substantial reduction in FDI inflows in the aftermath. For hydrological events, there is a huge fluctuation, likely because the effects of landslides and floods are substantially different from each other. One can see that the FDI inflows fluctuate substantially in the years following an event. Similarly to this, meteorological events also see little to no significant effect in the years following storm events, only for

climatological events one does find a persistent reduction in FDI inflows. A probable explanation is the FDI inflows into agriculture, especially palm oil, which are likely to decrease shortly following forest fires.

The results are too weak to generalize, but one can see that by using the same methodology it seems that especially geophysical events matter for Indonesian FDI.

A.7 Assessment of the Shock Metric

The entire methodology and interpretation of the results is strongly dependent on the way the shock is defined. In the binary treatment version of the shock parameter in the main paper, a district is considered treated if the economic center of the district is exposed to an earthquake with at least weak destruction potential and if people are recorded to be affected by the event. Figure A.12 illustrates the impact of changing

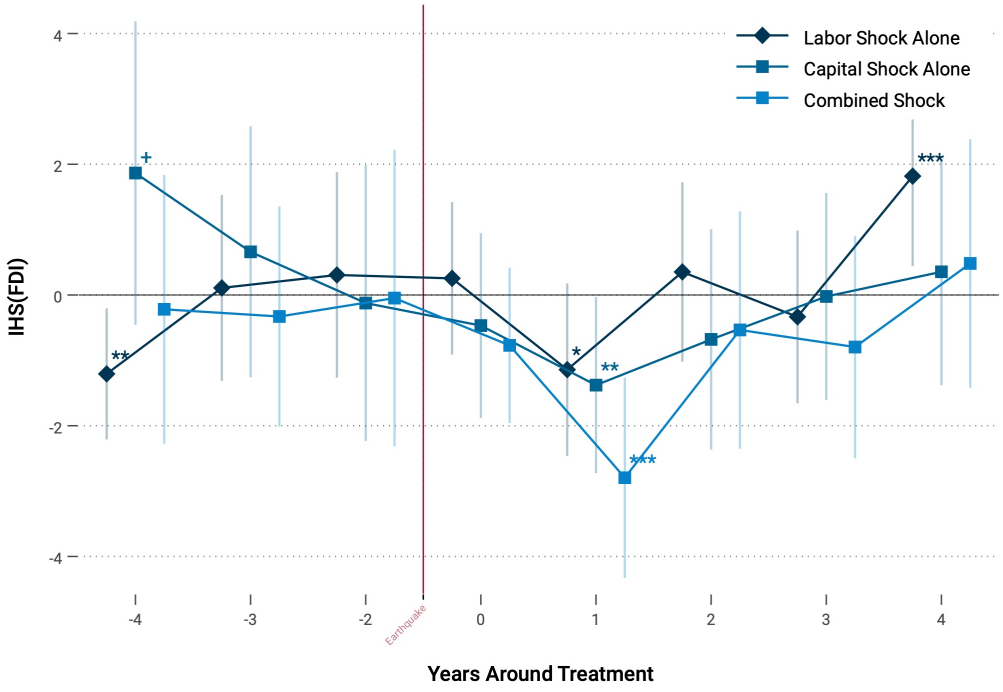

Fig. A.12. Definition of the Shock Variable

Note: The figure shows how the results of the baseline estimation changes if only the labor shock (light blue), capital shock (darker blue) or the combined shock, as used in the main paper (in darkest blue) is used. The corresponding treated sample size is 431 treated for labor shock, 207 for capital shock and 126 for joint shock. The confidence bands show the 95% interval with *** 1%, ** 5% and * 10% significance level.

the shock definition, namely using either the capital or the labor shock individually instead of jointly. A similar trend to that in the baseline can be observed. The larger coefficient in the combined shock, according to the definition used in this paper, can be attributed to a smaller sample size if both conditions hold. However, it also highlights that when labor and capital are affected simultaneously, the effects tend to be stronger.

Additional robustness tests include running the results based only on stronger earthquakes. The problem is that the sample size becomes so small, especially for extremely destructive earthquakes above a modified Mercalli scale level of 8 with 6 units, that no valid inference is possible without changing the estimation strategy. A similar case applies to thresholds for the minimum amount of local population that needs to be affected by an earthquake. In more than 90% of the cases, less than 5% of the local population was affected by losing their livelihood. Thus, the sample becomes incredibly small for valid inference. Below we provide additional robustness tests on the inclusion of intensity.

In addition, concern can be raised based on the correlation between the shock and the included baseline controls. All models include lagged log population estimates. However, the shock also includes at least a lagged population weight (i.e., the total number of affected persons as a share of the local population). Despite the shock being converted into a binary dummy, there might be some mechanical correlation between the shock and the lagged population. Furthermore, one might raise concerns of endogeneity if the treatment itself affects the population. To address both concerns, a sensitivity analysis is conducted that includes longer lags for the labor shock. Given that the shock today is unlikely to have influenced population two or three years ago, this is a valid check of the exogeneity of population. Additionally, at baseline, districts are restricted from having been exposed to previous earthquakes to avoid past treatments influencing the population. Using historical pre-shock data already reduces the likelihood that the shock itself impacts the population. Table A.7 shows the effects of the baseline estimation (however, only 3 lags are shown) by changing the weight for the labor shock. Instead of using the one-year lagged population variable, here we use the two-year (column (2)) and three-year (column (3)) lagged population variables to reduce the influence on today's values. The results remain significant.

Another test is to see if intensity and victim number really matter. The chosen model uses a binary treatment to define a shock and fit it into an event study framework. We assume a threshold of MMI VI in accordance with studies such as Gignoux and Menéndez (2016). To provide further evidence that this is not an arbitrary threshold,

Table A.7
Changing the Population Weight of the Labor Shock

	Labor Shock Only		
	(1) Pop (t-1)	(2) Pop (t-2)	(3) Pop (t-3)
Event Year	-0.224 [-0.982,0.534]	-0.231 [-0.990,0.528]	-0.197 [-0.941,0.546]
Event Year (Lag +1)	-0.962** [-1.866,-0.0568]	-1.004** [-1.937,-0.0711]	-0.972** [-1.923,-0.0205]
Event Year (Lag +2)	-1.028* [-2.200,0.143]	-1.078* [-2.279,0.123]	-1.082* [-2.317,0.153]
Obs.	4071.000	4071.000	4071.000
Baseline Controls	Yes	Yes	Yes
Time FE	✓	✓	✓
District FE	✓	✓	✓
Province× Year FE	✓	✓	✓

Note: *<0.1, **<0.05 and ***<0.01 p-value. Dependent variable: IHS(FDI inflows). Each model includes log NTL term. All models exclude treatment reversal and the simultaneous occurrence of other disasters. Districts that switched names are systematically excluded. Columns (1)-(3) change the definition of the weighting of the "labor shock" by including, instead of population lagged by one period (column (1)), also lagged by two and three periods.

additional tests are provided. To get a sense of how rare these events are, Figure A.13 shows the number of events that exceed level five by district, MMI and year. It shows that MMI VI (in darker blue) is found in every year in at least 8 districts. Stronger events, including VII and VIII, become increasingly smaller, and X is only found in 2004. Overall, this figure reveals that higher intensities indeed are increasingly more rare and that in every year we have at least one or two handful of districts that are exposed to hazards above the threshold. Next, it is assessed whether the results hold using the raw intensity and death numbers directly in the regression model. The toy regression to test intensity takes the following form:

$$f_{ijt} = \alpha + \gamma_1 \text{MMI}_{ijt} + \gamma_2 \text{Death}_{ijt} + \gamma_3 \text{Victim}_{ijt} + \Gamma X_{ij,t-1} + \omega_i + \omega_t + \omega_{jt} + \epsilon_{ijt}$$

This equation estimates the effect of an increase of one unit of MMI on inverse-hyperbolic sine transformed FDI. *Death* and *Victim* measure the number of people who died in that year from earthquakes, and victims are those who lost their homes. As in the baseline model, X includes the lagged log population and nightlight. The results are shown in Table A.8 in Column (1). Column (2) adds a lagged term to see if the results become

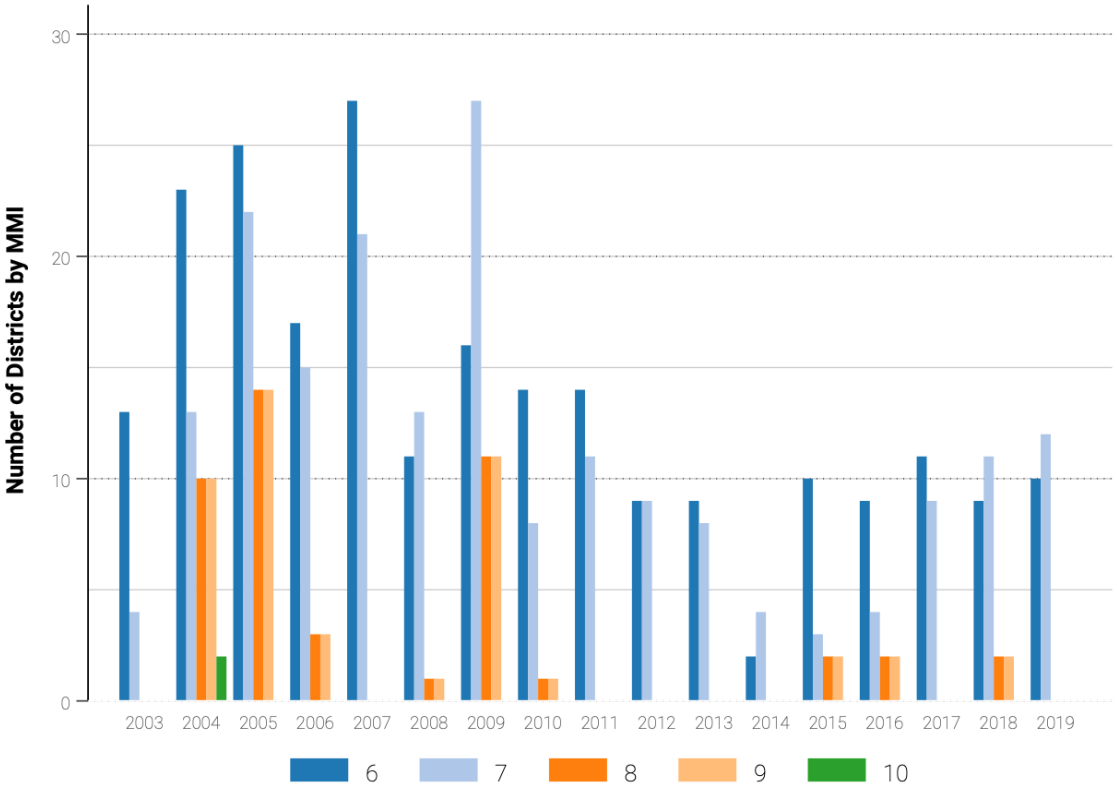

Fig. A.13. Distribution of Affected Districts (By Year and Intensity)

Note: The figure plots the number of districts each year that are affected by an earthquake above Modified Mercalli Intensity (MMI) V. In dark blue (MMI VI) are the "weakest" of these extreme events.

stronger in the year after, as shown in the main part of the paper. The table shows that MMI intensity has a strong and significant effect on the reduction of FDI inflows of around 8% per increase in intensity. This effect increases to 12% if lagged. The effects of death are more minor, but still significant. For each 100 people who die in this preliminary equation, around only 0.1% of FDI is reduced. Interestingly, victims, that is, people who lose their house, have a small positive effect.

This raises the urgent question of whether the intensity of MMI actually increases the number of people who die, victims, and, for robustness, also the houses reported as damaged in United Nations Office for Disaster Risk Reduction (UNDRR) (2024). The regression takes the form

$$f_{ijt} = \alpha + \gamma_1 \text{MMI}_{ijt} + \Gamma X_{ij,t-1} + \omega_i + \omega_t + \omega_{jt} + \epsilon_{ijt}$$

where f_{ijt} is either the number of people who died, the number of people who lost their

livelihood or the number of houses destroyed. The negative hyperbolic sine is used, given the excessive zeros in the outcomes. The results are shown in Figure A.14. When comparing the effects of no treatment with increasing intensity, it becomes clear that MMI V is a good threshold to define increasing levels of risk.

Overall, the results on changing the shock definitions and playing with the threshold, as well as the population dummy, reveal that the shock dummy itself potentially over-simplifies the nonlinearities of the hazard intensity. Increasing the threshold for more intense events causes the treatment sample size to become too small for inference. Future research may want to consider a more dynamic treatment model that does not rely on binary treatments and takes intensity into account more specifically.

Table A.8
Toy Equation: Allowing for Hazard Intensities

	(1) Asinh(FDI)	(2) Asinh(FDI)
MMI Intensity	-0.0747* [-0.151,0.002]	-0.0828** [-0.163,-0.002]
+100 Deaths by Earthquake	-0.00192*** [-0.003,-0.000]	-0.00201*** [-0.004,-0.000]
+100 Victims by Earthquake	0.0000*** [0.000,0.001]	0.0000*** [0.000,0.001]
MMI Intensity (Lag+1)		-0.1324*** [-0.213,-0.052]
+100 Deaths by Earthquake (Lag+1)		-0.00174*** [-0.003,-0.000]
+100 Victims by Earthquake (Lag+1)		-0.0000 [-0.001,0.000]
Obs.	5583.000	5258.000
Baseline Controls	Yes	Yes
Time FE	✓	✓
District FE	✓	✓
Province× Year FE	✓	✓

Note: *<0.1, **<0.05 and ***<0.01 p-value. Dependent variable: IHS(FDI inflows). Each model includes log lag NTL term and logged population. All models exclude treatment reversal and the simultaneous occurrence of other disasters. Districts that switched names are systematically excluded. Column (1)-(3) change the definition of the weighting of the "labor shock" by including instead of population lagged by one period (column (1)), also lagged by two and three periods.

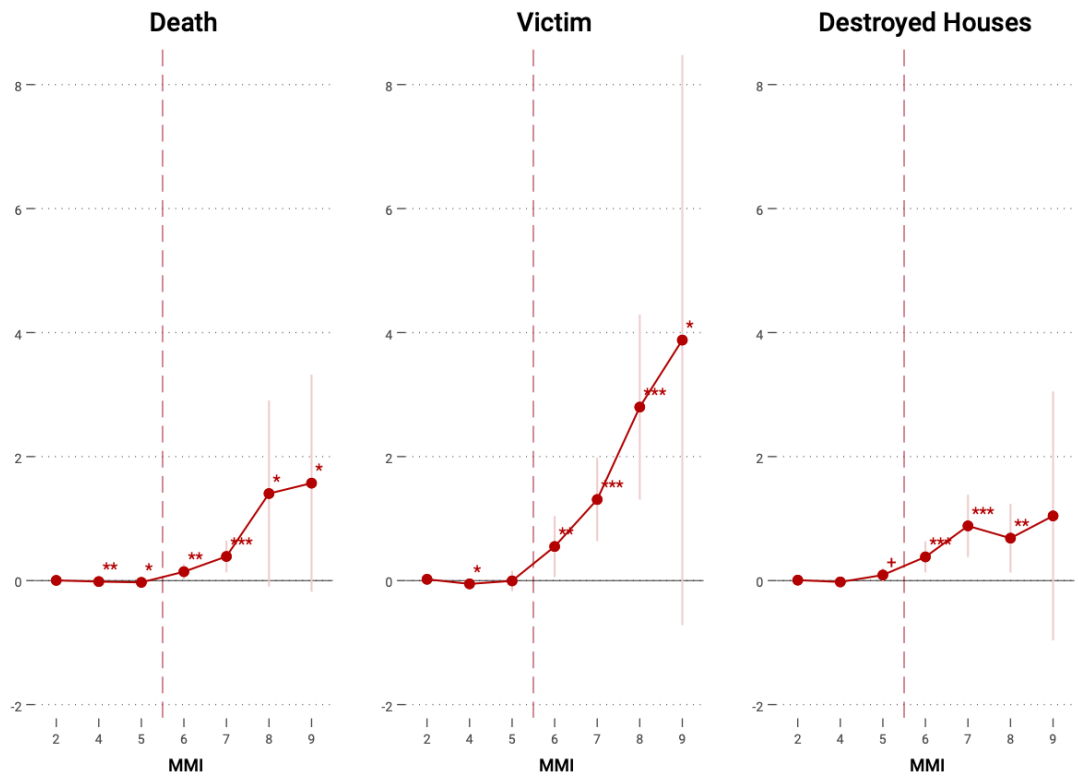

Fig. A.14. Exceeding the MMI V Boundary

Note: Exceeding the MMI V boundary. This figure plots in a simple toy equation if indeed earthquakes exceeding the MMI V threshold should be considered to be "destructive". The outcomes Death, Victims and Houses Destroyed are based on Desinventar United Nations Office for Disaster Risk Reduction (UNDRR) (2024). The models control for lagged population and nightlight and include district, year and province-year fixed effects. The dotted line represents the used threshold.

A.8 Robustness of Controls

Nighttime-light as Economic Proxy

Nighttime Light (NTL) is frequently used in the development literature as a proxy for economic (Sutton, Elvidge, and Ghosh, 2007; Hu and Yao, 2022). However, a key criticism is that it overvalues richer areas, whereas deprived regions are still insufficiently accounted for by nighttime light. This test serves solely to investigate the correlation between logarithmic nighttime light in Indonesia and logarithmic GDP at the provincial level. Data on GDP values[19] over time is drawn from the INDO-DAPOER data

[19] The GDP is based on 2008 constant values

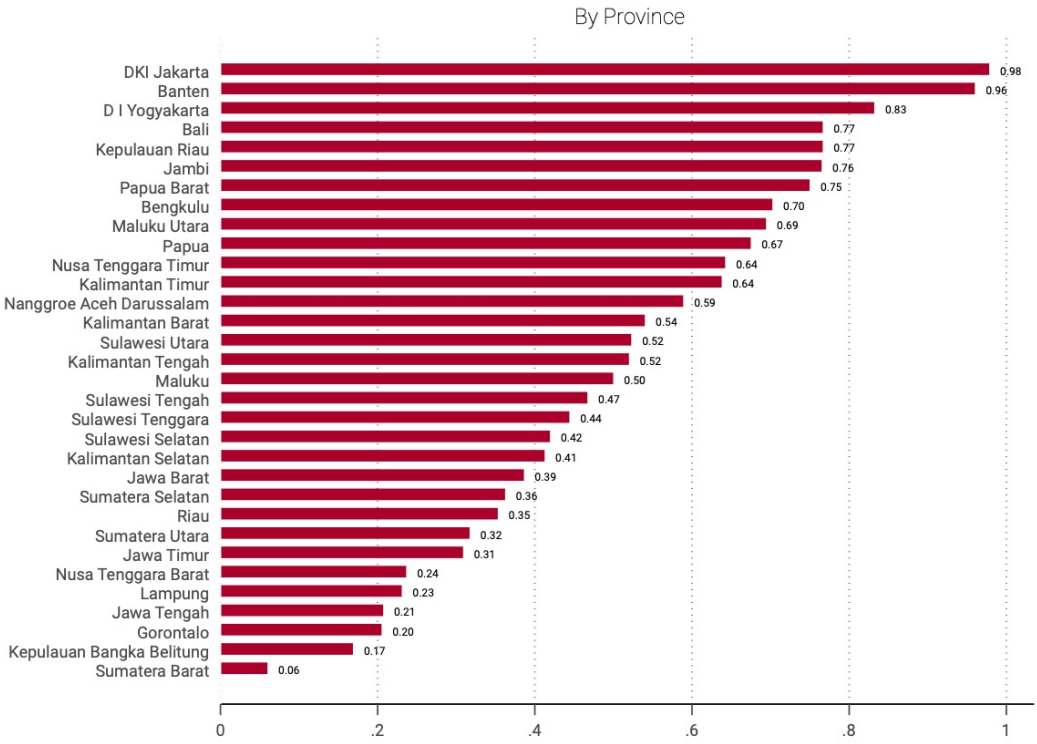

Fig. A.15. Correlation of Log(NTL) and Log(GDP)

Note: Data source from The World Bank (2022b) and Chen, Yu, et al. (2021). Author's own calculations.

frame of The World Bank (2022b). Figure A.15 shows the correlation by province over all years. The strongest correlations are found in DKI Jakarta and Banten, which are the richest provinces in GDP terms. Sumatera Barat and the islands of Bangka Belitung have extremely low correlations, despite ranking in the midfield for GDP ranking. The reason might be that there are many districts that are relatively poorly illuminated at night, and therefore the correlation between GDP and nighttime light is low for these districts. The mean correlation at the district level between NTL and GDP is 0.6, which is sufficiently high to suggest that nighttime light proxies GDP reasonably well for this dataset. On a yearly basis, the correlation is never lower than 0.58. So, despite the outlier of Sumatra Barat, NTL serves as a sufficient proxy for GDP.

Reliability of Economic Center Classification

Another important test is to assess if the economic centers within each district are accurately matched. To do so, information on the geographic coordinates of cities in Indonesia and their populations is obtained from GeoNames (n.d.). The cities are then

matched to the district level and overlaid with the Indonesian 5x5km grid. One has to keep in mind that economic centers are classified based on quintiles of nighttime light or population density. Using descriptive statistics, it is assessed, whether most of the included cities from GeoNames (n.d.) are actually located within the highest quintile in mean grid-cell specific nighttime-light. One can see in Table A.9 that most cities, indeed, fall into the fifth quintile. The cities in the first quintile are predominantly those that are later reclassified using the mean population density by grid cell. Then, by using information on the population within a city, it is also tested whether all major cities with more than a million inhabitants are classified in the highest quintile and indeed, all but one are. There is only one exception: the city of Depok. Analysis at the grid cell level revealed that the city boundaries extend over multiple grid cells, the neighboring grid cells fall into the highest quintile, yet the centroid of the city from the dataset does not. Thus, this outlier can be ignored, as Depok is taken care of by the neighboring grid cell.

Table A.9
Classification of NTL Quintiles onto City Geocoordinates

	Mean	Standard Dev.	N
First Quintile	14760.36	20868	14
Second Quintile	57184.5	80447	2
Third Quintile	14670.75	23722	4
Fourth Quintile	123394.2	227715	26
Fifth Quintile	163299.4	517996	364
Total	153729.1	492323	410

Note: Only grid cells which match city coordinates are included. Mean refers to mean population of each city in each quintile.

Reliability of Disaster Effect Data

A key criticism of the paper could be made if one does not believe in the main data source from United Nations Office for Disaster Risk Reduction (UNDRR) (2024). According to a report by UNDP Indonesia and BNPB Indonesia (2019) the management of the Desinventar database has been handed over from the UNDP to Indonesia's national agency for disaster management (*Badan Nasional Penanggulangan Bencana*) in 2008. It is continually improving in tracking and recording disaster occurrences and has also since launched its geodata portal (DiBI). In the following, we compared the Desinventar data to the two most common other data source in Indonesia: the offi-

cially reported data by the national agency BNPB (Badan Nasional Penanggulangan Bencana, 2022) and the data provided on event and aggregated district level by EM-DAT CRED and UC Louvain (2024).

Figure A.16(a) shows the percentage of district events matched between the two data sources. Differences in all categories, except forest fires and tsunamis, can be attributed to the fact that, before 2008, Desinventar had not been fully aligned with the national agency. Since 2008, there are no matching errors found. The difference in the tsunami originates from different accounting. Desinventar differs between tsunamis themselves and tsunamis with simultaneous earthquakes, whereas in the BNPB data, tsunamis and earthquakes would be coded as tsunami events. Forest fires are also for accounted by BNPB, but in the data portal, where this comparison data is obtained, we have not been able to download the category for forest fires as well. Nonetheless, on a national reporting level, these events are also reported. This test is shown to illustrate that Desinventar is not an arbitrary database, but the data is officially provided and managed by the National Agency for Disaster Management.

Moving on to the EM-DAT database, Figure A.16(b) indicates that Desinventar accounts for substantially more events than the EM-DAT database, whereas approximately 10% of EM-DAT events are not matched on a district level with the Desinventar data. In absolute numbers, only 136 events in total remain in EM-DAT out of which 15% are earthquake events, 22% are landslides, and 20% are volcano eruptions. Now, one must be cautious not to make overly optimistic claims based on that figure. First, it is possible that EM-DAT only accounts for additional districts for the same event, resulting in more or different districts being included. It is also possible that a certain percentage of matching errors result from discrepancies in district names. While it is attempted that all district names are standardized and collapsed by the 2003 Census IDs, it could nonetheless happen that some are mismatched. A key reason why Desinventar is likely to have substantially more events is that no threshold on when to include an event is used in contrast to the EM-DAT database. EM-DAT uses thresholds in which, for instance, at least ten or more people are killed or 100 or more are affected, whereas Desinventar includes any potential event.

A key advantage of using Desinventar over EM-DAT is the use of district- and event-specific damage and victim statistics. EM-DAT tends to report the total number of damage and/or victims on an event level. Even if it affects certain districts, the impact of the disaster may be unequally distributed across these districts and even within them. This discussion provides further validation of the usage and reliability of the

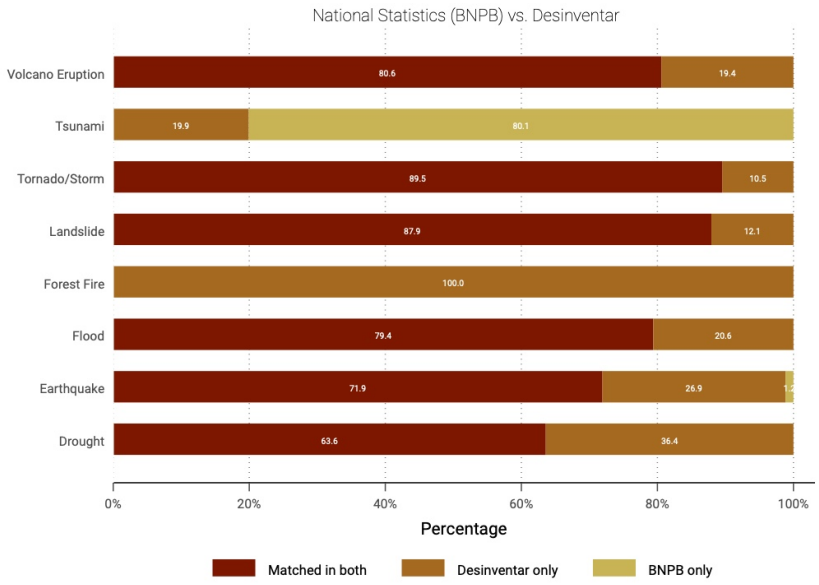

(a) Main Comparison

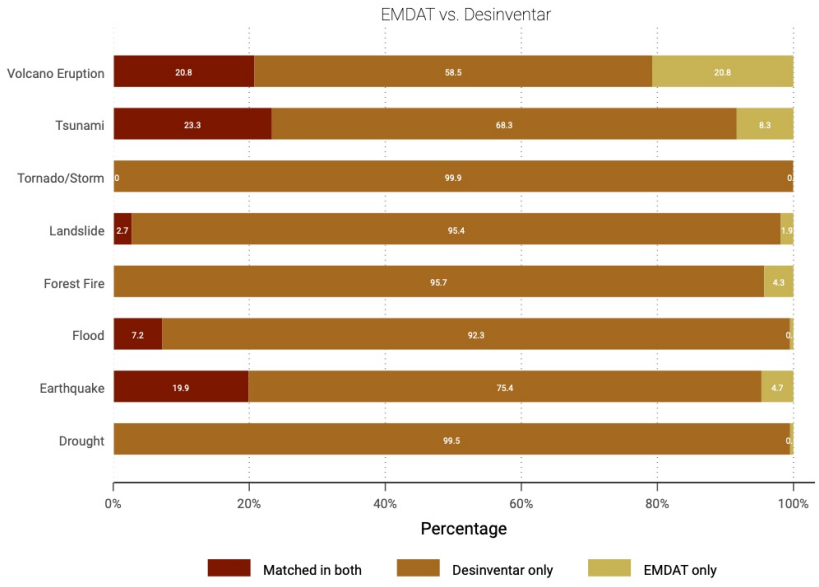

(b) Alternative

Fig. A.16. Comparison of Different Datasets

Note: Data is based on United Nations Office for Disaster Risk Reduction (UNDRR) (2024), Badan Nasional Penanggulangan Bencana (2022), and CRED and UC Louvain (2024) Calculations are the author's own. In both cases matching is done based on number of events recorded within districts of a year. The results should not be generalized and considered with caution as matching errors are likely to bias the 1:1 matching.

damage statistics used in the main part.

A.9 Assessment of Hazard Risk Indices

Creation

In order to account for risk, at first, models for disaster risk are obtained predominantly drawing from the Global Assessment Report on Disaster Risk Reduction of 2015[20]. These models show where different disasters are likely to take place based on historic data and probabilistic models. They serve the purpose of creating event-specific risk indices, which can be combined to an overall *composite hazard index*.

To create risk indices for the different disaster categories, it is essential to have a consistent measure of risk exposure. However, the five hazard maps used for different disasters (earthquake, flood, landslides, tsunami, and forest fires) have varying return periods, making direct hazard values incomparable. To address this, we merge the hazard maps with the 5x5 km grid of Indonesia, resulting in five grids, each representing a specific disaster type's risk per grid cell. These values are internally comparable within each disaster type but not across categories. To achieve cross-category comparability, we standardize the risk values using z-scoring. Each grid cell is then assigned a standardized risk value based on its respective category. District-specific values are obtained by averaging the z-scores of grid cells within each district, resulting in mean standardized risk scores for all Indonesian districts across the five disaster categories. For these five categories, the risk index is not endogenous to the occurrence of disasters, as it only represents the statistical likelihood of occurrence in a given year based on long-historical time series.
For volcanoes, storms, tidal waves, and droughts, where hazard maps are unavailable, we use a different method. Here their z-scores are based on the recorded number of events within each district, drawing from the Desinventar database (United Nations Office for Disaster Risk Reduction (UNDRR), 2024). This z-score represents the district's deviation in event occurrence from the average across all districts over 17 years,

[20] Flood Hazard 200 years (UNDRR, 2017a), Landslides - Frequency (triggered by Precipitations) (International Centre for Geohazards and Smebye, 2018), Landslides - Frequency (triggered by Earthquakes) (International Centre for Geohazards / NGI and Smebye, 2018), Tsunami - Frequency (International Centre for Geohazards / NGI and Lovholt, 2009), Fire Density (UNEP/DEWA/GRID-Europe and Peduzzi, 2014), Peak ground acceleration 250 years (PGA is a metric to measure earthquake intensity) (UNDRR, 2017b)

serving as a rough proxy for risk likelihood. Admittedly, the disaster risk for volcanoes, storms and droughts is endogenous to occurrences within the district, but might nonetheless reflect, to some extent the respective district's risk likelihood.

For an overall index of disaster-specific risk, a so-called *composite hazard index* is created. Following Emrich et al. (2022) first for each disaster category the averaged district-specific z-score is recoded to equal one if the score is below -0.5, two if the score is between -0.5 and 0.5, three if it is above 0.5 to 1.51 and four if it is above 1.51. These event-type specific dummies are then aggregated using the formula:

$$\text{CHI}_i = \sum_{c=1}^{n} x_{ic}$$

whereby x_{ic} refers to the recoded standardized risk score for district i and the disaster category c is. The composite hazard index reflects the district time invariant index for multi-hazard disaster risk, which will be used to validate if risk likelihood plays a role in investment response.

Assessing Validity of the Two Main Hazard Indices

The heterogeneity analysis in this work relies on the validity of the composite hazard index created and the dummy variable for earthquake risk to accurately reflect disaster risk in a region. To validate the accuracy of the computed indices, they are compared to existing studies.

Evaluating the Calculated Composite Hazard Index
In the first place, Figure A.17 shows the composite hazard scores computed by district. Note here that as the map is based on boundaries as in 2015. The risk indices, in contrast, as the entire dataset is based on the 2003 district boundaries, causing some districts to remain without data (in gray). They, nonetheless, have risk values assigned in the dataset, they are just not plotted. One can see high-risk scores (darker colors) around the main island of Java.

Evaluating the Calculated Earthquake Hazard
To evaluate the precision of the earthquake hazard dummy, it is visually compared to a pre-existing map from UN OCHA Regional Office for Asia Pacific (OCHA-ROAP) (2011), which is also provided in Gignoux and Menéndez (2016) in Figure A.1.

One can see a similar pattern of earthquake risk, if one confirms Figure A.18 with these

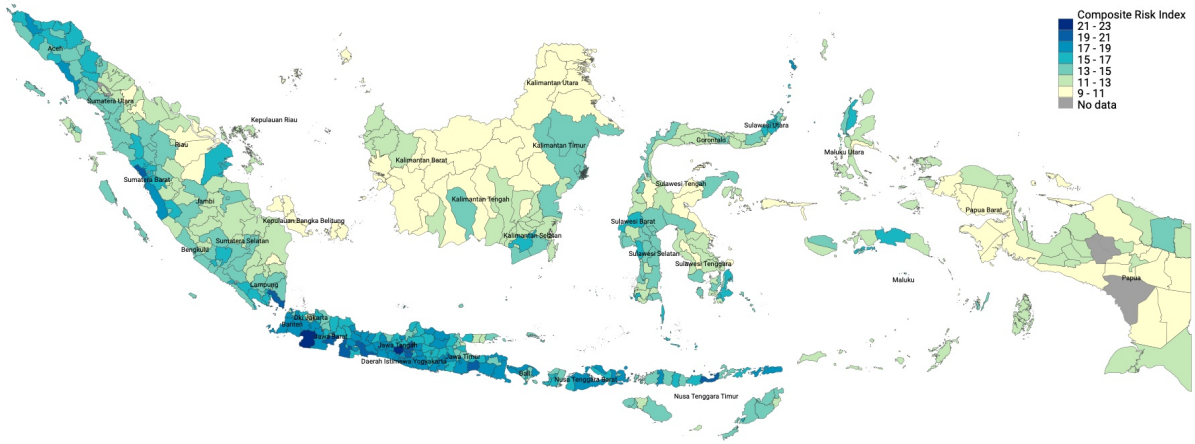

Fig. A.17. Visualization of the Hazard Index

Note: Shapefile is based on OCHA's Regional Office for Asia and the Pacific (2021). Composite Hazard Scores are self-computed scores based on probabilistic hazard maps as described in the respective section. Composition based on percentiles. In gray, there are the 14 districts that are not matched. The graph likely overestimates FDI slightly in 51 districts.

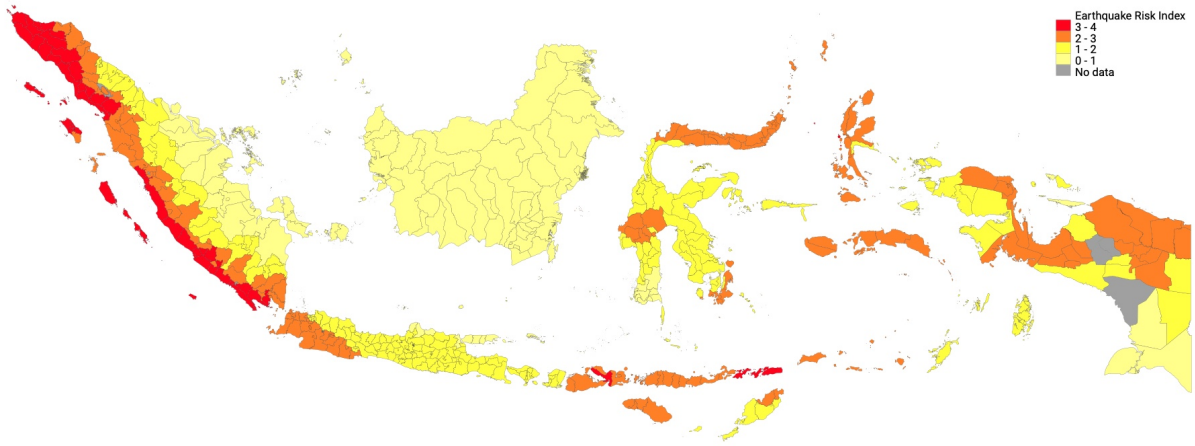

Fig. A.18. Visualization of the Earthquake Risk Metric

Note: Graph is based on own calculation, shapefile from OCHA's Regional Office for Asia and the Pacific (2021) and probabilistic hazard map based on 250-year PGA return periods (UNDRR, 2017b). 14 districts are not matched (in gray).

maps. This broadly confirms that the risk metric is broadly correct.

In order to see whether the results of the heterogeneity analysis change if another hazard metric is used, instead of a z-scoring based on the hazard maps, a z-scoring using the number of events is used, as mentioned. In comparison to the results of the heterogeneity analysis, the findings remain unchanged. It is also tested to see whether the results hold if the hazard scores are calculated using only the hazard within economic centers of each district. Again, the results do not change. Overall, it appears that despite the hazard score not being perfect, it proxies disaster risk in this setting sufficiently well to be used in the analysis.

CHAPTER 2

The Platform for Economic Analysis of Climate Hazards (PEACH): Gridded Hazards from 2000-2019

This chapter is solely authored.[1]

2.1 Background & Summary

As the likelihood of extreme events, including those that co-occur, increases, the need for higher-resolution analysis in impact studies becomes more critical. Improved resolution of such studies enables better tailoring of policies for future responses to extreme events, especially in the Global South. In the past, social scientists frequently estimated responses at the country level (Shabnam, 2014; Botzen, Deschenes, and Sanders, 2019), which underestimates local effects, at the household level (De and Thamarapani, 2022; Sen, Nayak, and Mohanty, 2022), often with gaps in years, or on specific hazards (Hsiang and Jina, 2014; Anderson et al., 2021), potentially underestimating the interaction between different types of hazards (Ridder et al., 2020). Most studies rely on pre-existing event-specific datasets such as the widely used Emergency Events Database (EM-DAT) (Rosvold and Buhaug, 2021a) or the DesInventar database (United Nations Office for Disaster Risk Reduction (UNDRR), 2024) hosted by the United Nations Office for Disaster Risk Reduction. These datasets are extremely valuable as they provide detailed information on the economic and societal effects of disasters. Yet, they are often provided at aggregated administrative levels (country or province level) and do not always report the intensity of an event; both are needed to better understand the interaction between society and the environment. Furthermore, EM-DAT has been criticized for overrepresenting richer countries (Felbermayr and Gröschl, 2014).

In response to upcoming criticism, there has been a shift to down-scale the data to a more local level in the GDIS project (Rosvold and Buhaug, 2021a) or to extend, notably

[1] I would like to thank Lisa Chauvet, David Daou, Alex Ruane, Greg Yetman, Maxime Souvignet, Andrew Kruczkiewicz, Manuel Linsenmeier, Kevin Schwarzwald, as well as participants of the Internal Reseach Seminar at Columbia University for their valuable feedback and discussions. All interpretations of the results remain solely my responsibility.

the EM-DAT database, with intensities for a subset of events in the so-called ifo-GAME database (Felbermayr and Gröschl, 2024; Felbermayr and Gröschl, 2014). The last major improvement in this type of literature has been the introduction of a new methodology to classify co-occurring, so-called "compound" events (Claassen et al., 2023a) using an algorithm called "MYRIAD-HES". The MYRIAD-HES project also provides data on hazard events for eleven hazards on a daily level between 2004 and 2017. The resolution of the hazard varies depending on the input data used. To determine which events to include, thresholds are established for each type of hazard. The data presented here is the closest comparable to the MYRIAD-HES data with three main differences.

First, the Platform for Economic Analysis for Climate Hazards (PEACH), the dataset described here, attempts to largely avoid imposing thresholds whenever possible. This is important as for many studies it is necessary to know also if an event is close to a threshold. For example, strong precipitation might be strong enough to damage crops, but not strong enough to be classified as an anomaly or an extreme event. The PEACH data includes, in total, 17 hazard categories on a monthly resolution on a global unified 10 km grid from 2004 to 2015. For 13 of the 17 hazards, data is available from 2000 to 2019. Thresholds are only applied when the input data is available at a higher temporal resolution (hourly or daily) and needs to be aggregated to a monthly value. This affects 6 of the 17 categories. Monthly resolutions come at the cost of not being able to examine daily co-occurrences as thoroughly as possible with the MYRIAD-HES event database (Claassen et al., 2023b). However, the approach chosen for PEACH allows researchers to set their own threshold of what defines an "extreme event," as this might vary depending on the study context. It also allows for a homogenized panel of hazard exposure over time.
Second, a key difference between PEACH and previous datasets is that this approach justifies the inclusion of hazards not by relying on the EM-DAT classifications of hazards. Instead, the hazards included are guided by the notion of "climate impact drivers" (CID) (Ruane et al., 2022; Seneviratne, Nicholls, et al., 2012), which are specific climate components that directly affect human systems and ecosystems. These CIDs are referenced by the Intergovernmental Panel on Climate Change (IPCC) and have been assessed for their potential to impact socio-economic and environmental outcomes. As they have already been evaluated, it helps researchers identify which hazards are most relevant to their specific context and where research gaps exist. Of the 17 hazards included, 14 are CIDs. The remaining three hazard classes are geophysical hazards that are not directly driven by climate change, but have substantial impacts on human

health and economic development. Taking into account 17 hazards allows us to further study the interconnectivity of hazards, which have been found to cluster around global hotspots (Lee, White, et al., 2024; Ridder et al., 2020).

Third and last, the PEACH data adds a selection of socio-economic variables to the dataset. While it is difficult to determine which variables are the most urgent, the chosen variables can be categorized into three classes, all of which are particularly relevant for social science studies: economic activity, hazard resilience/vulnerability, and agriculture.

For economic measures, nightlight is included. The inclusion follows the influential work of (Elvidge et al., 1997), which became widely applied in the economic literature through the work of Henderson, Storeygard, and Weil (2011). Economic activity due to electricity consumption appear as brighter spots in satellite images. Since this can originate from either economic activity or population, it's crucial to consider population size and density together with nightlight data. This is important for two reasons: first, rural areas with limited electricity access may have weak or no nightlight, where population estimates help fill data gaps. Second, nightlight can also reflect economic activity unrelated to human settlements, such as gas flaring or petroleum plants. Combining both variables provides a more comprehensive picture of human and economic activity, which is important for impact analysis.

To assess resilience and vulnerability, PEACH incorporates relative wealth and critical infrastructure. Advances in machine learning enable researchers to analyze large datasets —such as daylight images, building footprints, and more— to estimate wealth within a country. When studying the response to hazards, it is essential to consider the income level of a given area. While nightlight reflects economic activity, wealth is a broader measure of economic well-being. For example, a well-lit petroleum plant may not indicate local wealth, as nearby residents might be relatively poor due to lower housing costs. Similarly, the presence of critical infrastructure should be considered, as hazards affecting these areas can amplify negative consequences. Therefore, both relative wealth and critical infrastructure indices are included.

Finally, for agriculture variables, the Normalized Difference Vegetation Index (NDVI) as well as land use is added. The famous work of Reed et al. (1994) or Pettorelli et al. (2005) made the use of NDVI common in the literature. The satellite images used to measure vegetation allow for the study of changes at high temporal and spatial resolutions. In PEACH, due to its aggregation, some spatial detail is lost; however, trends over time remain useful for assessing large-scale changes in farming. However, within

each cell, only a portion may actually contain agricultural land. To address this, land use data is incorporated to weight NDVI values based on the fraction of vegetation-prone land. In addition, land use data complement population and nightlight estimates by offering insight into urban land distribution.

To conclude, the PEACH dataset effectively unifies well-established datasets with a globally consistent grid, providing easier access to data for researchers, particularly those in social sciences interested in environmental studies, disaster risk management, and adaptation to climate change. Methodologically, it primarily utilizes weighted statistics and combines data from various sources to create the dataset. Although not all CIDs are covered, the replication code is written to be easily expandable to additional CIDs. PEACH provides a global grid, yet for one of the 17 hazard classes, there is only quasi-global coverage. This concerns rainfall, as the data is derived from the high-resolution CHIRPS data (Funk, Peterson, Landsfeld, Pedreros, Verdin, Rowland, et al., 2014; Funk, Peterson, Landsfeld, Pedreros, Verdin, Shukla, et al., 2015) that reports precipitation only for the 50 ° N and 50 ° S latitudes in the current version. This is a known limitation, as it excludes snow and ice; nonetheless, CHIRPS is still used due to the high resolution of the data. All other hazard classes have global coverage. A subset of hazard classes is validated with ground truth data from weather stations from the Global Summary of the Day dataset (GSOD) (NOAA National Centers for Environmental Information, 1999) and compared with data from the EM-DAT data set. An overview of the data is provided in Figure 2.1. The upper panel (a) shows the exposure of individual cells with more than one person per square kilometer to the total number of hazards recorded by quantiles using a potential chosen threshold for extreme events. More information on the thresholds taken is provided in Appendix B.1. The lower panel (b) shows the covered CIDs and their corresponding impact dimensions, as taken from the IPCC report. In red, it shows the temporal coverage of the individual hazards included.

2.2 Methods

To create a unified framework for the PEACH data, the following steps are taken: First, a global grid is established. Then, the hazard classes are selected and transformed into the grid. Finally, socioeconomic data is added.

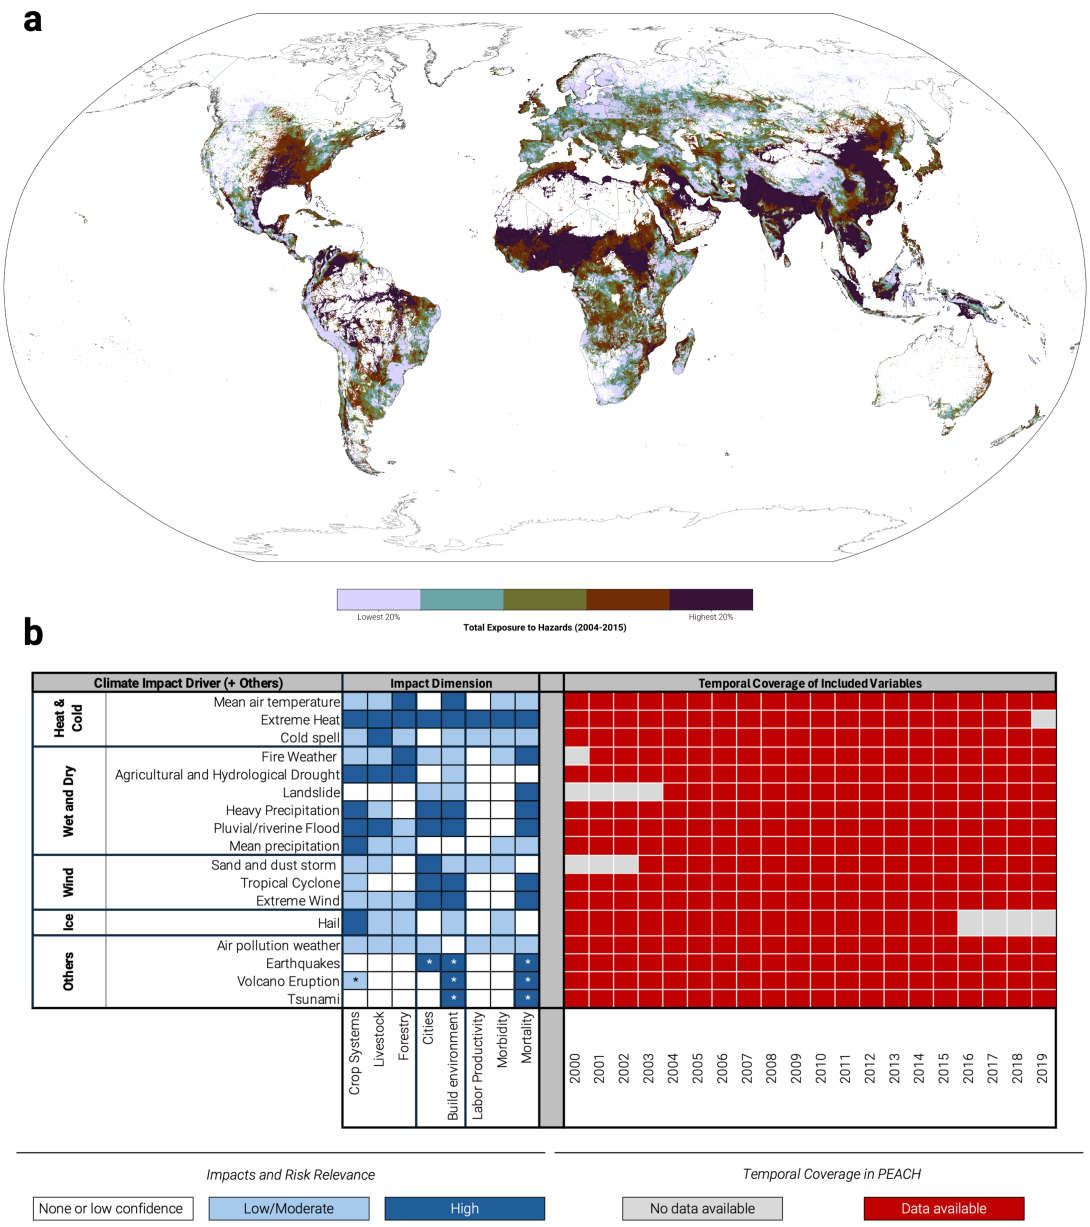

Fig. 2.1. Overview on the Data

Note: The figure shows the global distribution of hazard exposure, the hazard types included, their temporal coverage and their impact channels. Panel (a) shows the total sum of hazard exposure by cell for the years 2004 - 2015 for 15 considered types. For illustrative purposes, in this figure, cells with fewer than 1 person per square km are excluded. More information on the thresholds taken to define events is provided in the Appendix. Colors are defined based on scientific colors (Crameri, Shephard, and Heron, 2020). Panel (b) illustrates, on the left, the impact dimension on outcomes, adapted from IPCC Chapter 12, Table 12.2 (IPCC, 2021). The * symbol categories are not included in IPCC and are predicted by us. The right table plots the data coverage in PEACH for each hazard.

2.2.1 Grid Creation

At first, the administrative boundaries of countries are obtained from the geoBoundaries dataset (Runfola et al., 2020). GeoBoundaries provides a consistent and regularly updated version of administrative boundaries. The boundaries for 2023 are used. The boundaries are then divided into 10x10 km cells, which were created using the Pseudo-Mercator projection. While this leads to distortion effects towards the poles, the Pseudo-Mercator projection is chosen because it is widely applied. When grid cells overlap two country boundaries, the cell is assigned to the country with the majority of the cell area. The grid is then reprojected to the standard Mercator projection (EPSG: 4326) for further use. Each cell is assigned a unique cell identifier. To increase computational speed, the global grid is divided into subregions following the classification of the United Nations Statistics Division (UNSD, 2023). Disputed areas are assigned to the region in which they are located. The 10x10 km grid cell size represents a middle ground between the highest-resolution input data (approximately 90m) and the coarsest input data (up to 75 km). Most of the input data falls within the 5 to 25 km range, making 10 km a practical compromise that preserves spatial heterogeneity while keeping computational efficiency. This resolution is fine enough to capture meaningful subnational variation (e.g. local hazard exposure) yet coarse enough to ensure that most datasets can be meaningfully aggregated or interpolated without excessive data loss.

2.2.2 Hazard Classes

The PEACH data includes 14 climate impact drivers (CID) and 3 geophysically related hazards. Not all CIDs mentioned in the IPCC Sixth Assessment Report, Chapter 12 (IPCC, 2021), are included, mainly due to data availability constraints. Most snow and ice categories are missing, as well as all oceanic and coastal hazards. The three non-directly climate-driven hazards are included due to their potential to have substantial impacts on human health and economic development. An overview of all hazards, their sources, and resolutions is provided in Table 2.1. The merging procedure for combining the source data with the grid of each region is illustrated in Figure 2.2. For gridded source data, the process depends on whether the resolution of the source data is higher than or lower than the 10km of the grid cells. If the resolution is lower, then zonal statistics are computed based on the majority overlap between cells, relying on the Python *rasterstats* package. If it is higher, the computed statistics depend on the

hazard category. Whenever possible, the weighted mean is computed based on the *xagg* package (Schwarzwald and Geil, 2024) in Python. For polygon layers, such as earthquakes or floods, again majority overlap conditions are used, and the maximum intensity values are computed. For point geometries, the intersections are computed. In the following, for each hazard class, a discussion on the applied transformation and the final outcomes is included.

Mean air temperature & Cold Spells The monthly 2-meter mean temperature data from ERA-5 (Copernicus Climate Change Service, 2022) used here cover each month from 1960 to 2020 and are provided in Kelvin. The data is transformed into degrees Celsius and three standardized z-scores are calculated using the standard formula:

$$x_{z,m} = \frac{x_m - \bar{x_m}}{\sigma_{x_m}}$$

Where x is the monthly temperature in cell x at a time point t. $\bar{x_m}$ is the mean temperature of cell x and σ_x the standard deviation for a given month. The scores are always calculated in comparison to the previous months. The differences between the three z-scores are based on differences in the way the mean and standard deviation are calculated. The first one ("*temp_z_max*") uses all available months since 1960. This score provides the long-term anomalies of the temperature from historical monthly means. The second score ("*temp_z_30*") computes the mean and standard deviation based on a 30-year moving window, e.g. for January 2002 all January from 1972 to 2001 are used. This definition follows climatological standard normals, also applied by the World Meteorological Organization (WMO, 2024). This provides the medium-run anomalies of temperature. Lastly, the third score ("*temp_z_10*") is computed in the same way as the second score just with a 10-year moving window. It reflects the short-term anomaly. Using the three different definitions enables us to better investigate whether adaptation occurs and to distinguish between long-term climate shifts and short-term temperature variability. PEACH, therefore, includes four variables for the three z-scores and the monthly mean temperature in degrees Celsius. In addition to investigating mean air temperature, z-scores also allow one to compute monthly cold spells by taking negative anomalies on mean temperature (Felbermayr, Gröschl, et al., 2022), potentially in addition to a threshold on low mean temperature. This definition of cold spells is suitable for investigating the impact on agricultural productivity rather than on mortality.

Extreme Heat Extreme heat and extreme temperature are related concepts, but have dis-

Table 2.1
Overview on Hazard Data

Category	CID	Source	Metric Used	Aggregation Scheme	Spatial Resolution	Temporal Resolution	Temporal Coverage	Spatial Coverage
Heat & Cold	Mean air temperature	ERA-5	Degree (C°)	Anomaly	0.25	Monthly	1940-present	Global
	Extreme Heat	Mistry (2019)	WBGT	Days >Threshold	0.25	Daily	1970-2018	Global
	Cold spell	ERA-5	Degree (C°)	Anomaly	0.25	Monthly	1940-present	Global
Wet & Dry	Mean precipitation	CHIRPS	mm	Total	0.05	Monthly	1981-present	50°S-50°N
	Flood (river or pluvial)	DFO Rentschler, Salhab and Jafino (2022)	DFO Severity Cell affected by RP100 event	Maximum & Number Weighted mean of cell	P <0.01	Daily -	1985-present -	Global Global
	Heavy Precipitation	CHIRPS	mm	Anomaly	0.05	Monthly	1981-present	50°S-50°N
	Landslide	Global Landslide Catalogue + COOLR Global Fatal Landslide Database	Event count Event count	Total Total	P P	Daily Daily	2007-present 2004-2016	Global Global
	Drought (any type)	HYDRO-Jules	SPEI 3,6, 12	Weighted mean of cell	0.05	Monthly	1981-2022	Global
	Fire Weather	Fire Weather Index - ERA-Interim Fire Climate Change Initiative	Fire Weather Index Burned pixel (CI>0.9)	Days >Threshold Weighted mean of cell	0.7 0.003	Daily Monthly	1980-2018 2001-2022	Global Global
Wind	Severe Wind Storm	ERA-5 - CDS Toolbox iBTrACS	Wind gusts (m/s) Maximum wind speed (m/s)	Days >Threshold Maximum & Number	0.25 P	Hourly Hourly	1940-present 1850-present	Global Global
	Tropical Cyclone	Willis Research Network Global Tropical Cyclone Wind Footprint (WRN-TC)	Maximum wind speed (m/s)	Maximum & Number	P	Daily	1989-2020	Global
	Sand and Dust Environments	Copernicus Atmosphere Monitoring Service (CAMS) - EAC4	Dust Aerosol Optical Depth at 550nm	Days >Threshold	0.75	Hourly	2003-present	Global
		ERA-5 - CDS Toolbox	Wind gusts (m/s)	Days >15 m/s	0.25	Hourly	1940-present	Global
Snow & Ice	Hail	Prein and Holland (2018)	Hail Probability	Days >Threshold	0.7	Daily	1979-2015	Global
		ERA-5 - CDS Toolbox	CAPE (J/kg) and Convective Precipitation (mm)	Days >Threshold	0.25	Hourly	1940-present	Global
		Global Summary of the Day	Hail Days	Days	P	Daily	1973-present	Quasi-Global
Other	Air Pollution Weather	Shen et al. (2024)	PM 2.5	Weighted mean of cell	0.1	Monthly	1998-2022	Global
		Copernicus Atmosphere Monitoring Service (CAMS) - EAC4	PM 2.5 & PM 10	Weighted mean of cell	0.75	Monthly	2003-present	Global
	Earthquakes	Shakemap Catalogue	Modified Mercalli Intensity (MMI)	Maximum & Number	P	Daily	1900-present	Global
	Volcano Eruptions	Global Volcanism Program	Volcano Explosivity Index (VEI)	Maximum & Number	P	Daily	1345 BC - present	Global
	Tsunami	Global Historical Tsunami Database	Wave Height at Runup (m)	Maximum & Number	P	Hourly	2000BC - present	Global
		Pronk et al. (2024)	Coastal Elevation (m)	Weighted mean of cell	0.0003	-	-	Global

Note: The table shows the used CID, the measure, associated metric, how the data is aggregated, the spatial resolution of the origin data in degree (P stands for polygon), temporal resolution and coverage as well as spatial coverage. Sources: Mean air temperature: Copernicus Climate Change Service (2022) and Muñoz Sabater (2019), Extreme Heat: Mistry (2020), Cold spell: Copernicus Climate Change Service (2022), Mean precipitation: Climate Hazards Center (UCSB) (2023), Funk, Peterson, Landsfeld, Pedreros, Verdin, Shukla, et al. (2015), and Funk, Peterson, Landsfeld, Pedreros, Verdin, Rowland, et al. (2014), Flood: Brakenbridge and Dartmouth Flood Observatory (2023) and Rentschler, Salhab, and Jafino (2022), Heavy Precipitation: Climate Hazards Center (UCSB) (2023), Landslide: Kirschbaum, Adler, et al. (2010), Kirschbaum, Stanley, and Zhou (2015), and Froude and Petley (2018), Drought: Gebrechorkos et al. (2023), Fire Weather: Giuseppe et al. (2019), Vitolo et al. (2019), Copernicus Climate Change Service (2019), and Vermote and Wolfe (2015), Severe Wind Storm: Copernicus Climate Change Service (2024b), Tropical Cyclone: Gahtan et al. (2018), Knapp et al. (2010), Done (2022), and Done et al. (2020), Sand and Dust Environment: Inness et al. (2019) and Copernicus Climate Change Service (2024b), Hail: Prein and Holland (2018a), Prein and Holland (2018b), Copernicus Climate Change Service (2024b), and NOAA National Centers for Environmental Information (1999), Air Pollution: Shen et al. (2024) and Inness et al. (2019), Earthquakes: United States Geological Survey (USGS) (2017), Volcanos: Global Volcanism Program and Venzke (2023), Tsunami: National Geophysical Data Center (2024) and Pronk et al. (2024). *Note:*During the writing process, the ERA-5 CDS Toolbox that is used has been deprecated and been partially replaced (Copernicus Climate Change Service, 2024a). Otherwise, daily statistics can be computed using ERA-5 hourly data (Hersbach et al., 2023).

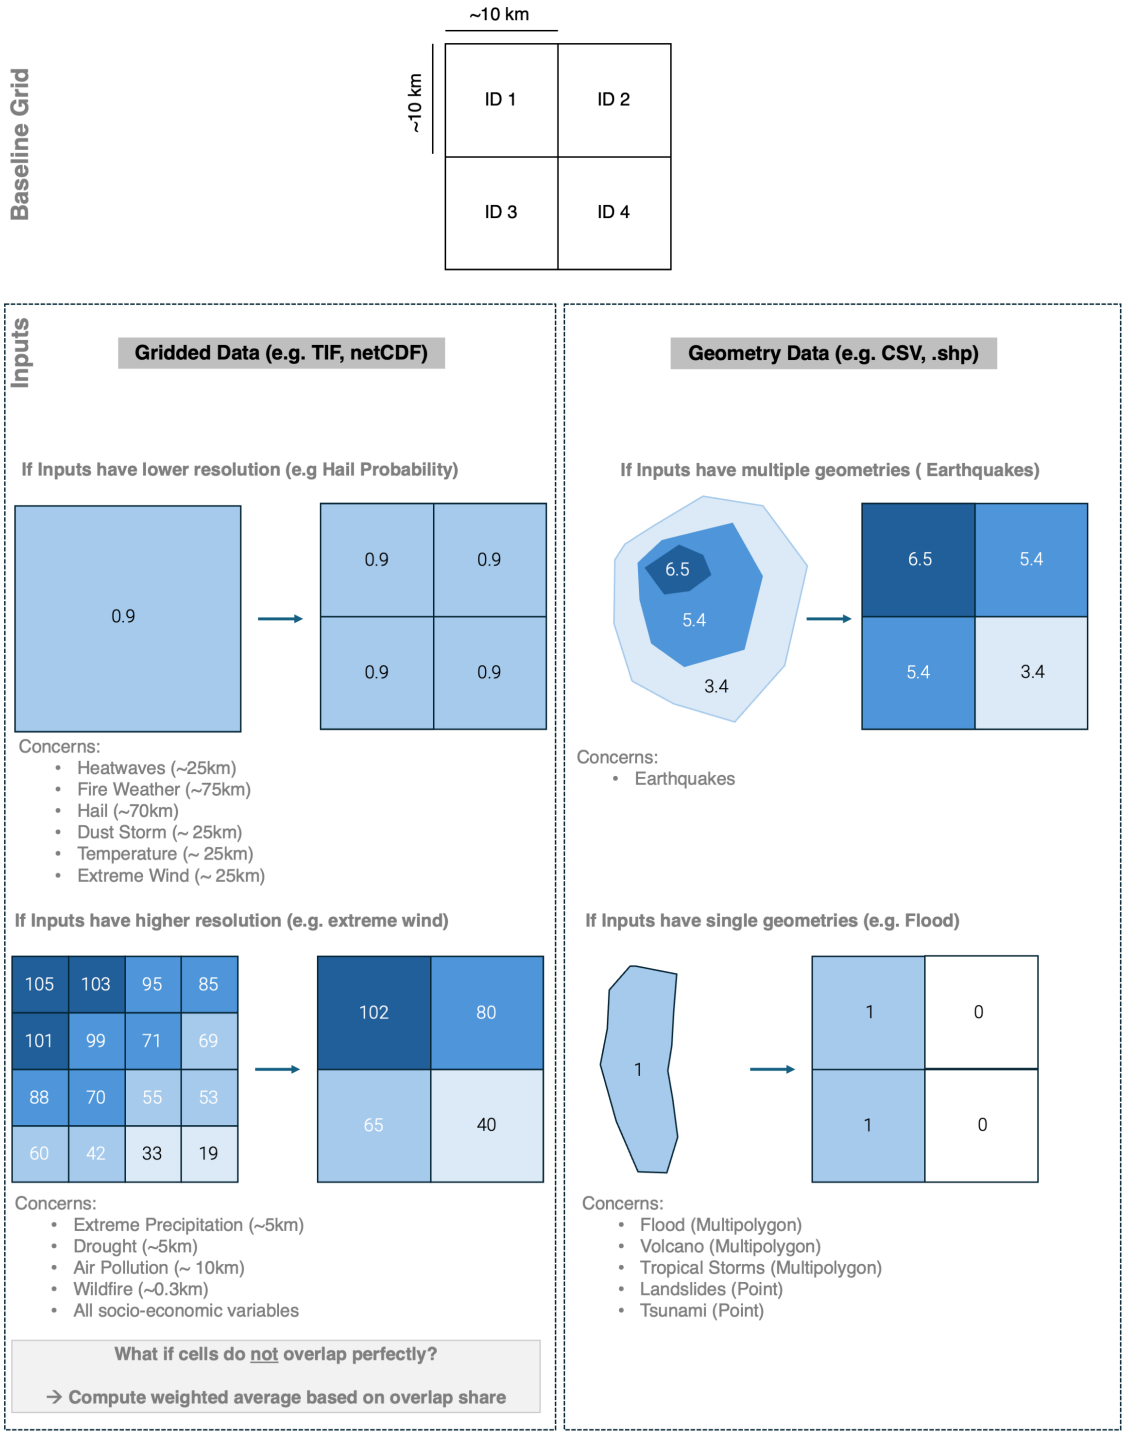

Fig. 2.2. Aggregation and Data Transformation

Note: This figure illustrates the spatial merging procedure to convert data from raster files or polygons onto a unified cell. The matching is conducted in Python and replication files are provided.

tinct meanings. In the context of PEACH, extreme heat is primarily concerned with its impact on human health, considering both temperature and humidity factors. Extreme heat is measured by daily maxima of the Wet-Bulb Globe Temperature (WBGT) - Index for the years 2000 to 2018 provided by Mistry (2020). To aggregate the data to a monthly level, two measures are calculated: maxima of rolling-means and the number of days above a fixed threshold. First, maxima of rolling means are computed as the mean WBGT for a one, three and five-day moving window. For example, if the WBGT is 30°C on day 1, 31°C on day 2 and 32°C on day 3, the rolling mean for a three-day window are 31°C. Then the corresponding values per month is the maximum of these one, three, and five day windows recorded. The second measure is the number of days when the WBGT exceeds a threshold. The threshold is set to 32.2°C, following the recommendations of the heat flag system (Patel, Mullen, and Santee, 2013; Clark, Konrad, and Grundstein, 2024), which is applied, for example, by the US military. Above this temperature, corresponding to a "black flag", physical exercise should be reduced to a minimum. The final data contains four final variables for the three different block maxima, as well as days above the threshold.

Mean Precipitation and Heavy Precipitation Information on heavy precipitation such as standardized score and total precipitation in mm per month comes from the CHIRPS database jointly with the Early Warning Explorer (EWX) (Climate Hazards Center (UCSB), 2023; Funk, Peterson, Landsfeld, Pedreros, Verdin, Shukla, et al., 2015). The data is available for every month, but unfortunately limited to 50°S-50°N, which needs to be kept in mind when trying to construct models for the utmost northern and southern regions (e.g. Siberia, Scandinavia, Northern Canada, Alaska and Patagonia). As the data is already provided in the correct format, it is merged with the grid using weighted zonal statistics. Two variables are included: the standardized rainfall anomaly score and the total precipitation.

Flood (River or Pluvial) The main input to flood is the boundaries of the affected flooded areas of the Dartmouth Flood Observatory (Brakenbridge and Dartmouth Flood Observatory, 2023). At first, for each event overlaps are computed with the global grid. For each event, the attributed given flood ID ("*DFO_ID*"), the number of deaths, displaced persons, severity, and duration are added. If multiple flood events are recorded in the same cell and month, the maximum values are taken. To control for the quantity, an additional dummy variable that counts the number of floods in each cell is added. However, flooded areas can be quite extensive. While not essential for flood analysis, recently published data (Rentschler, Salhab, and Jafino, 2022) on populations exposed

to flooding helps to pinpoint affected locations more accurately. The data is based on the Fathom flood model and measures the total number of population exposed to a flood with a 100-year return period for each 3 arc-second pixel. The exposure data is provided in tiles (areas of interest) with multiple bands, where each band contains various layers for different levels of inundation depth. First, it is assessed if all tiles overlap with the global grid. Following the authors' recommendations, all bands with inundation depths above 15 cm are combined to identify a population at high risk. Then, if at least one person is at risk in the pixel, the pixel is considered exposed. Aggregation with the grid is done by computing the average share of pixels at risk in each cell ("*flood_risk*"). To obtain a likely location of exposure to a flood, it is recommended to combine the flood risk variable with the flood observation record. The final dataset contains a variable for the flood ID, maximum number of deaths, displaced persons, severity, duration, count and the respective (socio-economic) cell risk.

Landslides Landslides are fairly local events and there are not many global datasets. To get a good coverage of events, the two most common datasets are combined: the Global Landslide Catalog (GLC) (Kirschbaum, Adler, et al., 2010; Kirschbaum, Stanley, and Zhou, 2015; NASA, 2023) and the Global Fatal Landslide Database (Froude and Petley, 2018). The reason is that the temporal coverage between the GLC and GFLD varies, which is why most events before 2007 are drawn from the GFLD and those after 2016 are exclusively from the GLC. The GLC is a project started in 2007 by NASA and is complemented in this version with the extended user-driven Cooperative Open Online Landslide Repository (COOLR). It provides point coordinates of landslide records, along with accuracy, possible trigger source, fatalities recorded, and date. Although active updating began in 2007, many events from the period prior have been added ex-post. PEACH includes landslides that have a specified location and date. Of the remaining 35 000 events, those are excluded where the accuracy is not known or is less than 25 kilometers accurate. This excludes another 22 000 events. Then to avoid double accounting, the same criteria are applied to the second dataset used. The GFLD accounts only for landslides in which fatalities are recorded. Following a similar approach than a recent study (Gómez, García, and Aristizábal, 2023) the two datasets are merged based on date and location. To avoid mismatches based on location precision, locations are rounded to the nearest kilometer. In total, 78 duplicate events are excluded. From the remaining events five variables are generated: first, if the trigger for the event is specified, a dummy is generated indicating the trigger. Then, depending on the accuracy, if it is below 25km, a buffer is drawn around the

point coordinate indicating the potentially exposed area of the event. To identify these events, a dummy is generated indicating the precision of the event location. Then, the coordinates of the points are merged with the grid. If multiple events are recorded in a cell for the same month, a counter variable is added to indicate the number of events in that month. If there are multiple event records, the trigger dummy and precision dummy are adjusted accordingly. If there are no multiple events, then the corresponding event identifier is included. Additionally, a dummy variable is added to indicate whether the event is drawn from the GLC-COOLR or the GFDL. The final data to add to PEACH accounts for the landslide identifier, the trigger cause, the precision of the event and a dummy if the landslide caused any human death and accounts for 16 800 events.

Drought The used HYDRO-Jules (Gebrechorkos et al., 2023) already reports monthly Standardized Precipitation Evapotranspiration Index (SPEI) on a five-kilometer resolution. To account for different types of drought, the durations are included: the 3-month SPEI, the 6-month SPEI and the 12-month SPEI. The 3-month SPEI evaluates short-term meteorological drought, the 6-month SPEI extends it rather to a medium-term meteorological drought; both indicate, however, also to some extent agricultural drought. The 12-month SPEI assesses rather long-term drought conditions indicating issues of general aridity. The PEACH dataset therefore contains three included drought measures.

Fire Weather and Wildfires Information on fire is divided into fire risk and fire exposure. Data on fire risk is based on the daily Fire Weather Index (FWI) (Giuseppe et al., 2019), which has a daily resolution of 70 km. To compute a monthly estimate, two metrics are computed. First, the monthly 5-day moving average maxima, so the highest value within a 5-day moving window period. Second, using all years from 1980 until 1999 to build the mean and standard deviation, a z-score is computed for each month. To measure exposure to fire, information on burned area (Copernicus Climate Change Service, 2019; Vermote and Wolfe, 2015) on a 300m per pixel is derived based on MODIS data provided and hosted by the ESA Climate Change Initiative and the Fire CCI. First, to increase computational speed, the resolution is coarsened to one kilometer, so that each aggregated pixel is the share of land burned in this one square kilometer area. Only pixels classified as burnt with at least 90% probability are considered. Then, for each global grid cell, the share of land burned is computed. For very large land masses, the merging part is compartmentalized into smaller subsets. A second variable is created that measures the burned land in areas identified as cropland (irrigated, rainfed, and

mosaic land). The final added data contains, therefore, four variables: the FWI 5-day block-maxima, the FWI anomaly, the share of the cell burned, and the share of burned land that is identified as cropland.

Severe Wind Storm The computation of severe windstorms is comparable to the MYRIAD-HES approach (Claassen et al., 2023b). The daily maximum of hourly wind gusts at 10m in m*s^{-1} from the CDS-Toolbox for ERA-5 (Copernicus Climate Change Service, 2024b) is acquired. Aggregating short-term wind gusts on a monthly level requires some decisions. Two statistics are computed: the monthly maximum of hourly wind gusts to see the strongest events and the number of days when wind gusts above 50.4 knots (26 m/s) are recorded, which is defined by the US Federal Emergency Management Agency (FEMA) and National Oceanic and Atmospheric Administration (NOAA) as minimum criterion for severe thunderstorm environments (Zuzak et al., 2023). Two additional variables are introduced: the maximum wind gusts reported in a month and the number of days with wind speeds exceeding 26 meters per second.

Tropical Cyclone Including tropical storms requires the computation of wind speed models around the center of the cyclone. As for the MYRIAD-HES (Claassen et al., 2023b) data, cyclone tracks from iBTraCS (Gahtan et al., 2018) are combined with wind field buffers from the Willis Research Network Global Tropical Cyclone Wind Footprint dataset (WRN-TC) (Done, 2022) based on the storm season and name. If no wind field buffer is available from WRN-TC for a specific storm, a 20-kilometer buffer is drawn around the eye of the tropical cyclone, and for each movement point of the storm track, the maximum wind speed in meters per second is included. Preferably, the NOAA-NHC one-minute averaged wind speed is used; if not available, the next largest estimate from the any agency affiliated to the World Meteorological Organization is included. For each cell and month, the maximum wind speed recorded is computed. Should multiple cyclones strike the same cell in the same month, the maximum recorded wind speed is taken, and a counter variable counts the number of storms within the same month. PEACH, therefore, contains two variables: the maximum wind speed of an event and the number of cyclones recorded in a given month.

Sand and Dust Environments To the best of our knowledge, there is no current dataset for dust storm events, yet the presence of high dust paired with stronger wind gusts could indicate at least dusty environments. To include a measure of days with high dust presence, as an indicative proxy of dust storm presence, two sources are combined: first, the 6-hour dimensionless Dust aerosol optical (DAO) depth at 550 nm (Inness et al., 2019) is converted into daily maxima. Then the 90^{th} percentile is computed to

capture only extreme dust environments. A dust environment day is defined if the daily maximum is above the 90^{th} percentile in optical depth and at least 0.5. Although this is a low threshold, similar thresholds have been applied in research (Logothetis et al., 2021). In desert-like areas, the percentile will ensure a more extreme event. Second, we add the hourly wind gust data from ERA-5 (Copernicus Climate Change Service, 2024b). A threshold of daily maximums of around 15 m/s is set. As we rely on wind gust and not sustained wind speed a slightly higher threshold than 10 to 14 meters per second as found in some studies (Muhammad Akhlaq, Sheltami, and Mouftah, 2012) is used. The two data sets are homogenized using spatial interpolation. A dusty day is defined as a day when both conditions hold: high levels of dust presence and daily wind gusts above 15 m/s. The included variable counts the number of days in a month in which both conditions hold. When used in a study, this variable should be taken with caution. It is indicative of days with high dust presence combined with stronger winds. Sand and dust storms are notoriously hard to model, so these variables could also be indicative of a blowing dust event, but not a full-blown dust storm.

Hail Data on global hail occurrence is scarce. Here, we primarily rely on a compiled daily database of hail probabilities (Prein and Holland, 2018a) from 1970 to 2015, which has been utilized in similar studies (Ridder et al., 2020). The resolution is with 0.75 degrees (or around 80km) quite large, but to our knowledge, it is the only available dataset applicable worldwide. The hail probabilities are estimated based on calibrated ground truth data from the US, Australia and Europe given the availability of the data. To create meaningful hazard variables, we compute two indicators: the number of days in a month that exceed a 90% and 99% hail probability. Both serve as lower-bound estimates of the events. For some parts of the world, this data set lacks ground truth data. Meteorological conditions might make hail probable, but the hail might melt in high-temperature areas on the way to Earth. This should be taken into account when using the variable. To further add information on likely hail locations, a metric of thunderstorm environments that could cause hail is created. To achieve this, based on reported hail events from weather stations, global thresholds are defined for two indicators used to assess the stability of thunderstorm environments. If then both indicators exceed the fitted thresholds defined by weather station data, an indicator variable is assigned the value 1, 0 otherwise. Finally, the number of days when the indicator equals one is counted. Therefore, from a global sample of weather stations (NOAA National Centers for Environmental Information, 1999), all reported hail events between 2000 and 2020 are extracted. Then, for each location, the corresponding daily maximum hourly

values for the two chosen indicators, Convective Available Potential Energy (CAPE) and Convective Precipitation (CP) (Copernicus Climate Change Service, 2024b), are added. CAPE and CP are amongst other factors indicators for unstable thunderstorm conditions (Taszarek et al., 2021). The 93 687 hail records of 7 417 stations are binned into latitude bins of 1 degree for CAPE and 5 degrees for CP. Next, within each bin the 95^{th} percentile of CAPE within that bin is computed as well as the median value of CP. The median is chosen as excessive rainfall might hinder hail formation. Using the computed percentile or median in each bin, a polynomial function is fitted (4th order for CAPE and cubic for CP) to get a continuous function of thresholds by latitude for CAPE and CP. Mathematically this corresponds to:

$$\widehat{\text{CAPE}}(\phi) = \alpha_4\phi^4 + \alpha_3\phi^3 + \alpha_2\phi^2 + \alpha\phi + \alpha_0 \tag{2.1}$$

where ϕ is the observed CAPE_{95} in a latitude bin. A similar approach is used for CP:

$$\widehat{\text{CP}}(\phi) = \beta_3\phi^3 + \beta_2\phi^2 + \beta\phi + \beta_0 \tag{2.2}$$

where ϕ corresponds to the CAPE_{50} in a latitude bin. The fitting is shown in Figure 2.3. To avoid inclusion of too weak convection events CAPE is limited to 150 J*kg^{-1} and CP to a minimum of 0.25 mm*h^{-1} as defined for thunderstorm environments in literature (Taszarek et al., 2021). Most thresholds for latitude bins are beyond 1000 J/kg CAPE, which has also been associated with hail formation for a study applied to Europe (Kaltenböck, Diendorfer, and Dotzek, 2009). Using hourly data aggregated to daily maxima (Copernicus Climate Change Service, 2024b), days in which observed values exceed the CAPE and CP latitude-specific thresholds are considered to be days with thunderstorm environments potentially favorable for hail formation. For PEACH, the number of days each month is then counted. In the future, this indicative measure should be improved by using alternative hail indices, such as the Significant Hail Parameter (SHIP), or at least by adding information on the 0-6 km vector wind shear. The final dataset will have three variables: the number of days when hail has been likely and the number of days with thunderstorm environments potentially favorable for hail formation.

Air pollution weather To account for ambient air pollution two sources are utilized: monthly mean fine particulate matter at 2.5 μm ($\text{PM}_{2.5}$) for the full period (Shen et al., 2024) and, at a lower resolution, data from an alternative data source (Inness et al., 2019) on monthly mean particulate matter at 10 μm (PM_{10}) from 2003 to 2019. The

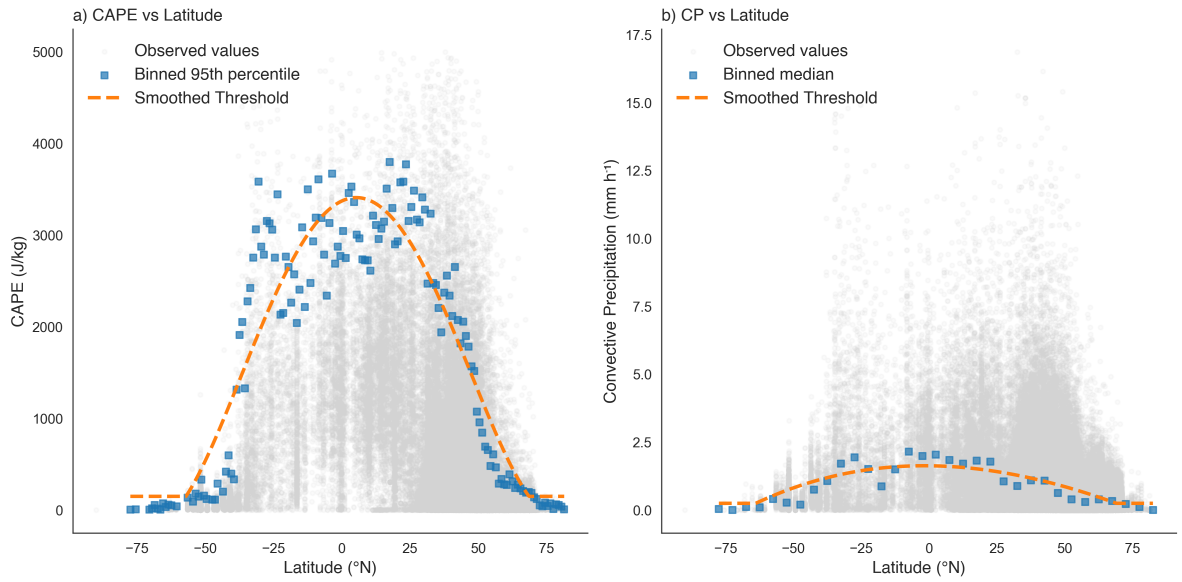

Fig. 2.3. Smoothing and Fitting the Hail Thresholds

Note: The figure shows the smoothing of the thresholds for Convective Available Potential Energy (CAPE) and convective precipitation (CP) by latitude. The thresholds are based on the 95^{th} percentile of CAPE and the median of CP in each latitude bin. The thresholds are smoothed using a polynomial function. The gray dots represent reported hail events at weather stations. For illustrative purposes, 79 hail events with a CAPE>5000 are excluded from plotting.

data comes in different units. The $PM_{2.5}$ is provided in $\mu g * m^{-3}$, while PM_{10} originally is provided in kg*m^{-3}. Thus, PM_{10} is converted into $\mu g * m^{-3}$. PEACH includes three variables: two $PM_{2.5}$ estimates from both data sources, and one for PM_{10} from the second.

Earthquakes The intensities of earthquakes are taken from the USGS ShakeMaps product (United States Geological Survey (USGS), 2023), which provides an interpolated intensity on an event basis for most recorded earthquakes. Two datasets are obtained from USGS using the Command Line Interface: All events with a minimum moment magnitude of 4.5 recorded and all that include polygons of intensities (so-called "shake maps"). First, for all epicenters reported, a 400 km buffer is drawn to see if they intersect with any cell or are uniquely located in the ocean. Then, for every reported ShakeMap, the average and highest Modified Mercalli Intensity (MMI) in every cell is computed. For all events where no ShakeMap is available, the grid is merged with the remaining epicenters, and the maximum reported magnitude is included. Magnitude here can refer to various types of magnitude (e.g., moment, regional, or body-wave), but is taken preferentially from the official USGS-reported magnitude of the event.

The number of earthquakes in a cell is counted, and if only a single event occurred, the USGS identifier is included. PEACH therefore contains five variables: whenever possible the USGS identifier of the event, the total number of events, if ShakeMaps are available, the mean and maximum reported MMI, and for non-ShakeMap events the maximum magnitude.

Volcano Eruption Having obtained information on all reported volcano eruptions (Global Volcanism Program and Venzke, 2023), the magnitude of the eruption is included based on the *Volcanic explosivity index* (VEI). If no end date is reported, the eruption is excluded. That concerns 33 eruptions. Following the MYRIAD-HES methodology (Claassen et al., 2023a), the radius of influence of a volcano is calculated based on the formula:

$$L = 3.0408e^{0.6956VEI}$$

The final variables included are the VEI of the eruption in the area of influence, as well as the volcano number.

Tsunami From the Global Historical Tsunami Database (National Geophysical Data Center, 2024) the subset of Tsunami run-ups for the period covered is obtained. The analysis is predominantly based on the arrival time and the maximum water heights in meters. Although the run-up data also reports the horizontal inundation depth for some events, the issue is that proximate areas with similar elevation might also be affected but without having the location recorded. To also include such grid cells, for analysis a digital elevation model is included for all coastal cells (Pronk et al., 2024). The resolution is on a 30m. To combine it with the 10km grid, the median elevation of each cell is calculated by first scaling down the 30m resolution to 1km and then computing the median elevation. Around each reported run-up location, a 25 km buffer is created, which serves as lower bound estimation. The idea is to account for at least one neighboring cell of a treated cell. Cells with reported tsunamis are always considered treated, while nearby cells are considered treated if the observed tsunami wave height is greater than the median elevation of the cell. For PEACH, the included information is the maximum wave height in a month recorded, as well as, if possible, the corresponding NOAA Tsunami identifier.

2.2.3 Socioeconomic Data

PEACH encompasses a range of commonly used variables to assess the economic and social performance of an area. The list of included variables is based on the availability of higher-resolution socioeconomic outcomes.

Nighttime Light In many developing countries local measures of economic activity is scarce. Nighttime light has commonly been applied to measure the effects on economic activity (see for instance recent discussions by Felbermayr, Gröschl, et al., 2022; Pérez-Sindín, Chen, and Prishchepov, 2021; Hu and Yao, 2022). Data on annual nighttime light is based on a homogenized time series between the two standard sources: the Defense Meteorological Satellite Program (DMSP) and the Visible Infrared Imaging Radiometer Suite (VIIRS) (Li, Zhou, et al., 2023). The data is on a 1km resolution and for each cell globally the mean, median and maximum nighttime light intensity are computed. Following the authors' advice, cells with a light intensity below seven are considered to be 0 to avoid the inclusion of stray light.

Population Population is an essential factor to measure the potential impact of a hazard event. The data is taken from the annual GlobPop dataset (Liu, Cao, et al., 2024a) and is available with a resolution of 1km. The data is provided in two variables: the total number of people and the population density per square kilometer. Both are obtained and merged with the grid: the population is aggregated at a 10km resolution as the total sum of population counts within a cell that lies within its boundaries. The population density is calculated as the weighted mean within the cell.

Relative Wealth To measure the relative wealth of an area, the microestimates of wealth from Chi et al. (2022) are included. The relative wealth of a 5km cell is computed by combining household surveys, satellite images, cellular network data, and other anonymized data from Facebook. Using all these data sources, different types of classification and deep learning algorithms are used. The final relative wealth index provides information on whether an area is, in relative terms, wealthier or less wealthy compared to other parts of the country. The variable is time-invariant and aggregated to the 10km grid by using the minimum, mean and maximum recorded in each cell as well as the aggregated mean error. The data is available for 135 middle and low-income countries.

Critical Infrastructure To measure the resilience of an area, information on critical infrastructure is included. The data is taken from Nirandjan et al. (2022) and is time-invariant. The data is derived from OpenStreetMaps of November 2020 and is pro-

vided on a 0.1-degree resolution. The source data provides an index from 0-1 on the global scale. To compute the index subindices for telecommunication, health, education, waste, energy transportation, and water (so-called "*systems*") are calculated based on initial OpenStreetMap assets. For studies investigating the effects of climate impact drivers or hazards in general, the type of infrastructure might matter, the mean critical infrastructure spatial index (CISI) as well as the score for every system is included. To homogenize with PEACH the weighted sum within each cell of the index is computed using the overlapping area share between the CISI index and the 10km cells.

Normalized Vegetation Index For many agricultural settings, it is crucial to know the plant health of a region. The Normalized Difference Vegetation Index (NDVI) is a standard measure for assessing plant productivity. Information on monthly NDVI is taken from a combined homogenized time series of the Advanced Very High Resolution Radiometer (AVHRR) and Moderate Resolution Imaging Spectroradiometer (MODIS) (Li, Cao, Zhu, et al., 2023) since 1982 on 8km resolution. Following the author's recommendation, values with scarce vegetation below the value 0.1 are excluded. To obtain a valuable measure for PEACH, three variables are included and calculated: first, the weighted monthly mean NDVI in the cell, which serves as a general measure of plant health. Then based on the last 10 years and 18 years a moving standardized z-score is computed for each month. The moving z-score shall ensure that changing climate conditions are taken into account. PEACH thus contains three variables: mean NDVI, NDVI anomaly based on monthly expected values in the last 10 years and the previous 18 years.

Land Use Lastly, to calculate the annual changes in land use over time, which are essential for adaptation studies, information on 1km land use is taken from the Historic Land Dynamics Assessment (HILDA +) project (Winkler et al., 2020). As land classes for each cell the share of urban land, cropland, pasture or range land, forest, unmanaged grass or shrub land and sparse land with no vegetation is computed.

2.3 Data Records

The Platform for Economic Analysis of Climate Hazards (PEACH) is a derived dataset for 14 continental climate impact drivers (CID) and 3 other hazards with data on socioeconomic, infrastructure and environmental data. It aims to provide easier accessible access to disaster and economic data. The data is on a global grid on 10km resolu-

Table 2.2
Overview on Socioeconomic and Environmental Data

Category	Name	Source	Metric Used	Aggregation Scheme	Spatial Resolution	Temporal Resolution	Temporal Coverage	Spatial Coverage
Socioeconomic Factors	Nighttime Light	Li et al. (2020)	Light Intensity	Weighted mean of cell	0.01	Yearly	1992-2021	Global
	Population	Liu et al. (2024)	Population Count & Density	Weighted mean of cell	0.01	Yearly	1990-2020	Global
	Relative Wealth	Chi et al. (2022)	Relative Wealth Index	Weighted mean of cell	0.05	-	2022	Quasi-Global
Infrastructure	Critical Infrastructure	Nirandjan et al. (2022)	Critical Infrastructure Index	Weighted mean of cell	0.1	-	-	Global
Environmental	Normalized Vegetation Index	Li et al. (2023)	NDVI Anomaly	Mean and Anomaly	0.08	Monthly	1981-present	Global
	Land-Use	Winkler et al. (2020)	Land Use Share	Mean	0.1	Yearly	1960-2019	Global

Note: The table shows the included socio-economic measures, the measure, associated metric, how the data is aggregated, the spatial resolution of the origin data in degree, temporal resolution and coverage as well as spatial coverage. Sources: Nightlight: Li, Zhou, et al. (2023) and Li, Zhou, et al. (2020), Population: Liu, Cao, et al. (2024a) and Liu, Cao, et al. (2024b), Relative Wealth: Chi et al. (2022), Critical Infrastructure: Nirandjan et al. (2021) and Nirandjan et al. (2022), NDVI: Li, Cao, Zhu, et al. (2023) and Li, Cao, Zaichun Zhu, et al. (2023), Land Use: Winkler et al. (2020) and Winkler et al. (2021).

tion projected in EPSG:4326. An exception is rainfall that is only quasi-global up to 50° latitude. It has full coverage for the years 2004 to 2015 with more than 80% of all variables being available between 2000 and 2019. Data is provided in the Apache Arrow Feather format that is usable in most software applications. Sample code is provided to help you get started in both R and Python should you not be familiar with either.

PEACH is openly accessible on Zenodo upon acceptance. The data is provided in sub-regions of the World based on the United Nations Statistics Division (M49) classification. The data is due to licenses of input data under the Creative Commons Attribution-NonCommercial (CC BY-NC) license.

The data used is compiled from already peer-reviewed published articles, meaning the validity can be assured. Errors would occur only during spatial merging with the grid. To assess accuracy data, two tests are conducted: first, temperature, precipitation, hail, and wind speed is validated against ground-truth data from the Global Surface Summary of the Day (GSOD) (NOAA National Centers for Environmental Information, 1999). Second, the commonly used EM-DAT (CRED and UC Louvain, 2024) and its downscaled version (Rosvold and Buhaug, 2021b) are matched with the PEACH data to see if certain event types or regions are underidentified.

Validation based on weather station data

At first, data from 8 000 to 12 000 weather stations, depending on the year, for daily reported meteorological statistics is taken. The locations of each weather station overlap with the PEACH grid and aggregate on a monthly level, similar to how PEACH is generated. The same metrics for wind are calculated as they are computed in PEACH. Three different biases are evaluated: monthly deviations, yearly deviations and geographic deviations. It is assessed if values vary trend-wise by month, year and region. The results are shown in Figure 2.4.

The temperature is, on average, underestimated by half a degree (0.25 degrees median) in comparison to weather station data on a monthly basis. The standard deviation is around 1.78 degrees. The largest bias is around April. Regionally, most parts of the world do not see large deviations; only some weather stations in western China, Turkiye, the west coast of the United States, and the Pacific coast of Latin America see larger deviations. On an annual level it appears that the difference is slowly decreasing over time.

For precipitation, it becomes visually clear that PEACH is limited to latitudes of 50 °. The dataset sees on average 3.8 mm (1.2 mm median) in total precipitation, less rain than the individual dataset from the weather station, with a standard deviation

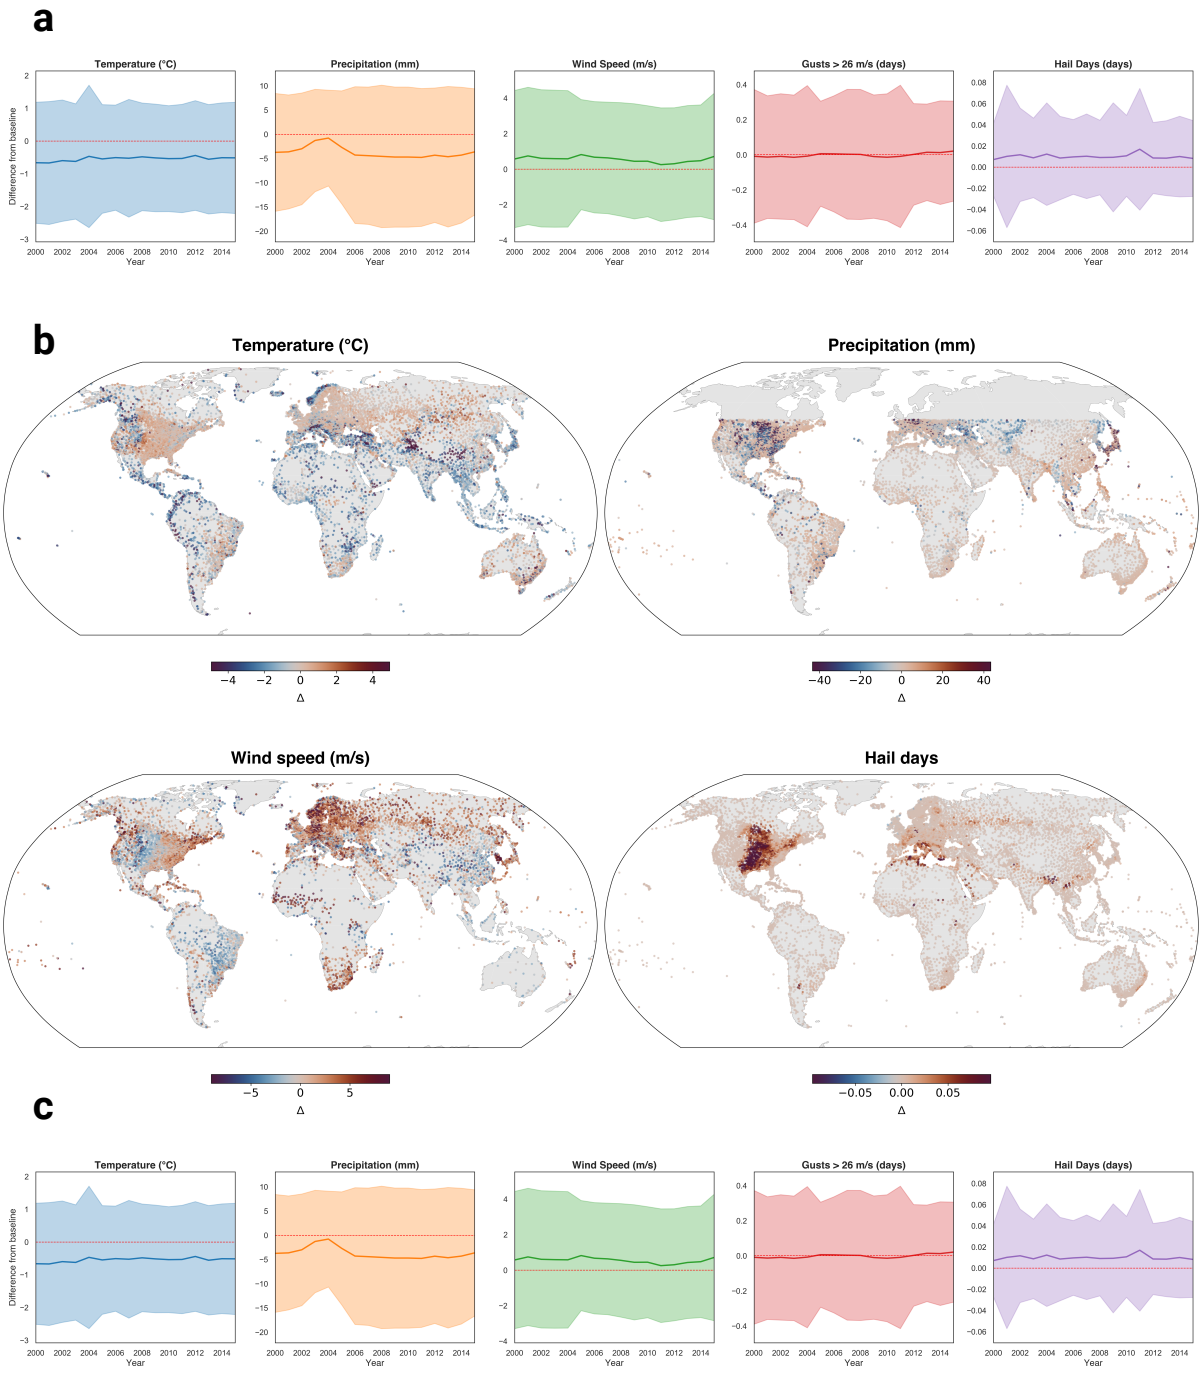

Fig. 2.4. Validation of PEACH data with Global Summary of the Day

Note: Panel (a) Shows the mean difference between PEACH and GSOD by month. In dark blue is the mean difference and in lighter blue the standard deviation. Panel (b) shows the mean difference by variable and station. Blue color indicates an underestimation of the value by PEACH and red color an overestimation. Panel (c) shows the mean difference by variable and year. In dark blue is the mean difference and in lighter blue the standard deviation.

of around 13 mm. The largest bias occurs in December and January, particularly in the Northern latitudes, where the bias becomes more pronounced. This might be due to precipitation taking place especially in Northern hemisphere winter months more often in the form of snow, as well as lower light availability in Winter, which might be affecting satellite-based infra-red measures. The bias does not decrease over time.
For monthly maximum wind gusts, PEACH overestimates the wind speed by 0.62 m/s (0.5 m/s median) with a standard deviation of 3.5m/s. There is no clear annual trend, but the summer months in the Northern Hemisphere see rather underestimation than overestimation. The days of wind gusts above 26 m/s are broadly comparable with weather station data and see on average less than 0.001 days of difference. The largest bias for wind geographically arises from Northern Europe. whereas parts of western China, Brazil, and the Midwest of the United States rather see an underestimate of wind gusts.
Finally, hail days for days with a 99% likelihood of hail are on average 0.01 (0 days median) day overestimated to station records. The low overestimation is due to the training of the hail probability data also drawn from weather station data. The standard deviation is around 0.05 days. No clear trend monthly or annual is observable.

Validation based on EM-DAT

In order to compare EM-DAT data with PEACH some data transformation is needed. First, EM-DAT (CRED and UC Louvain, 2024) and its scaled-down version of Geocoded Disasters (GDIS) (Rosvold and Buhaug, 2021b) are downloaded and information on hazard types and dates are added to GDIS. If for a unique disaster identifier (*disasterno*) an entry in GDIS is included, it is not taken from EM-DAT and information on type, start month, year and country as well as the geometry of the affected administrative boundaries are included. Otherwise, the information is taken from EM-DAT at the country level. The data is then aggregated to the 10km grid. The data is then compared with PEACH. The results are shown in Figure 2.5. Based on the disaster subtypes classification only relevant hazard classes are included (exclusion of not-included types in PEACH like *Tornado*, *Avalanches* or *Epidemics* or types where the associated class is unclear e.g. *severe weather*).
In the next step, the associated PEACH data for each country is loaded and prepared. To compare it with EM-DAT, thresholds are applied to the intensities to classify them as a hazard event. For flooding, the flood severity reported by the Dartmouth Flood Observatory (Brakenbridge and Dartmouth Flood Observatory, 2023) must be reported. For an earthquake, either the magnitude or the modified Mercalli intensity needs to

be above four. For hail, at least one day in the month is reported to have hail with a probability greater than 90%. For heatwaves, the three-day block maximum of the Wet Bulb Global Temperature needs to be above 28 degrees. For a storm event, the monthly maximum in wind gusts must be greater than 15 meters per second. For wildfires, at least one percent of the cell is burned in the corresponding month. For dust storms, at least one day of dusty environment days is counted. For tropical storms, inclusion is based only on wind speed. For landslides, the count of landslides in the month is taken. Cold waves are defined as a monthly deviation from the historic 30-year mean of more than one standard deviation. Droughts are considered if SPEI-3 is less than -1.3. For a tsunami, the tsunami count is taken. And finally, for volcano eruptions, a positive Volcano Explosivity Index (VEI) is used.
Lastly, the three data sets are merged based on the following criteria: for GDIS an event is considered a match, if the threshold for the same hazard type is exceeded for any cell of PEACH that falls within the affected administrative boundaries as reported in GDIS in the same start month and year. If either the threshold is not exceeded, or the month and year or the event does not fall into any of the boundaries, the event is considered as not reported in PEACH. For EM-DAT, the same is reported; however, instead of using administrative boundaries, the full country is considered. Only events between 2003 and 2015 where there is full coverage in PEACH is included.

Figure 2.5 shows the results of the matching process by hazard type, database, and geographic region. On the left side, one sees the total share of matched events where for each disaster event a match is found. If multiple regions are affected by an event, as long as a match is found in any of the affected area, the event is considered matched. There is substantial heterogeneity in the matching. Panel(a) shows that the results by hazard type differed by the GDIS and the EM-DAT sample. For the GDIS sample, the highest matching is achieved in tsunamis (100%), earthquakes (93-100%) and volcano eruptions (89-97%). The lowest matching shares are found for dust storms (33%), hail (36%) and flood(48-69%). For dust storms, the low percentage is likely due to a small sample size (3 events) and may also be an underestimation of dust storm events. Given that per definition, dust storms are defined if the 90^{th} percentile is exceeded, the low matching might be due to the days only capturing the most extreme dust events in dusty regions. The same applies for hail events, as the potential counts of days with a probability of 90% are too conservative and do not account for events below this threshold. The lower matching for floods is worrying, yet explicable. EM-DAT utilizes multiple sources, including the Dartmouth Flood Observatory (DFO) data, which is

used to report flood events in this database. The usage of different sources to define flood events might cause the DFO data to be more often misaligned with the EM-DAT data set. However, the widespread use of the DFO data and their higher accuracy of affected regions motivate the usage of the data. Heatwaves, landslides, and windstorms all achieve a matching around 80%, where mismatches are spatial mismatches for landslides or different definitions of storms for heat waves and windstorms. Cyclone mismatches result mostly from either the country not being primarily affected by the cyclone impact (e.g., impact through flooding) or by a mismatch of months. In panel (b), the matching by country is shown. In about 36% of the countries more than 80% of the events are matched, and a similar share is observed for the above 60% matching share. The lowest matching is found in Burundi (25 events and 16% matched), Mongolia, Libya and Gabon (each one event). In addition, the figure shows that the majority of countries have less than 50 events recorded, which makes the sample size for matching country-by-country relatively small. Countries with large land surface and a high number of reported events such as China, India or the United States all perform above the 80% matching quality. If mismatches occur, a likely reason can be found in the misalignment of the starting month or the non-reporting in the used input data.

Although the matching quality varies, the validation overall shows that PEACH is a relatively reliable dataset for the set of climate impact drivers tested. The dataset is especially useful for middle and low-income countries where national data is scarce. One might use more accurate national data sets when working on Western Europe, China, Australia, and the United States. The validation also revealed that it tends to underestimate the presence of dust storms, hail events, and floods, while for the other hazard categories, it can be connected to the vast majority of EM-DAT events.

2.4 Usage Notes

The PEACH data can be used as a global, regional or national dataset, and the files are provided by the ISO-3 classifier stored in subregions of the world. The definition of ISO-3 and the subregions follows the classification of the United Nations Statistics Division (M49). The dataset is provided in the Apache Arrow Feather format, which can be read in Python and R with the *feather* package. In the download folder there is also an easy-to-use Python and R script to convert the data into a CSV format and to plot the data. Although it is possible to convert the files into a Microsoft Excel format

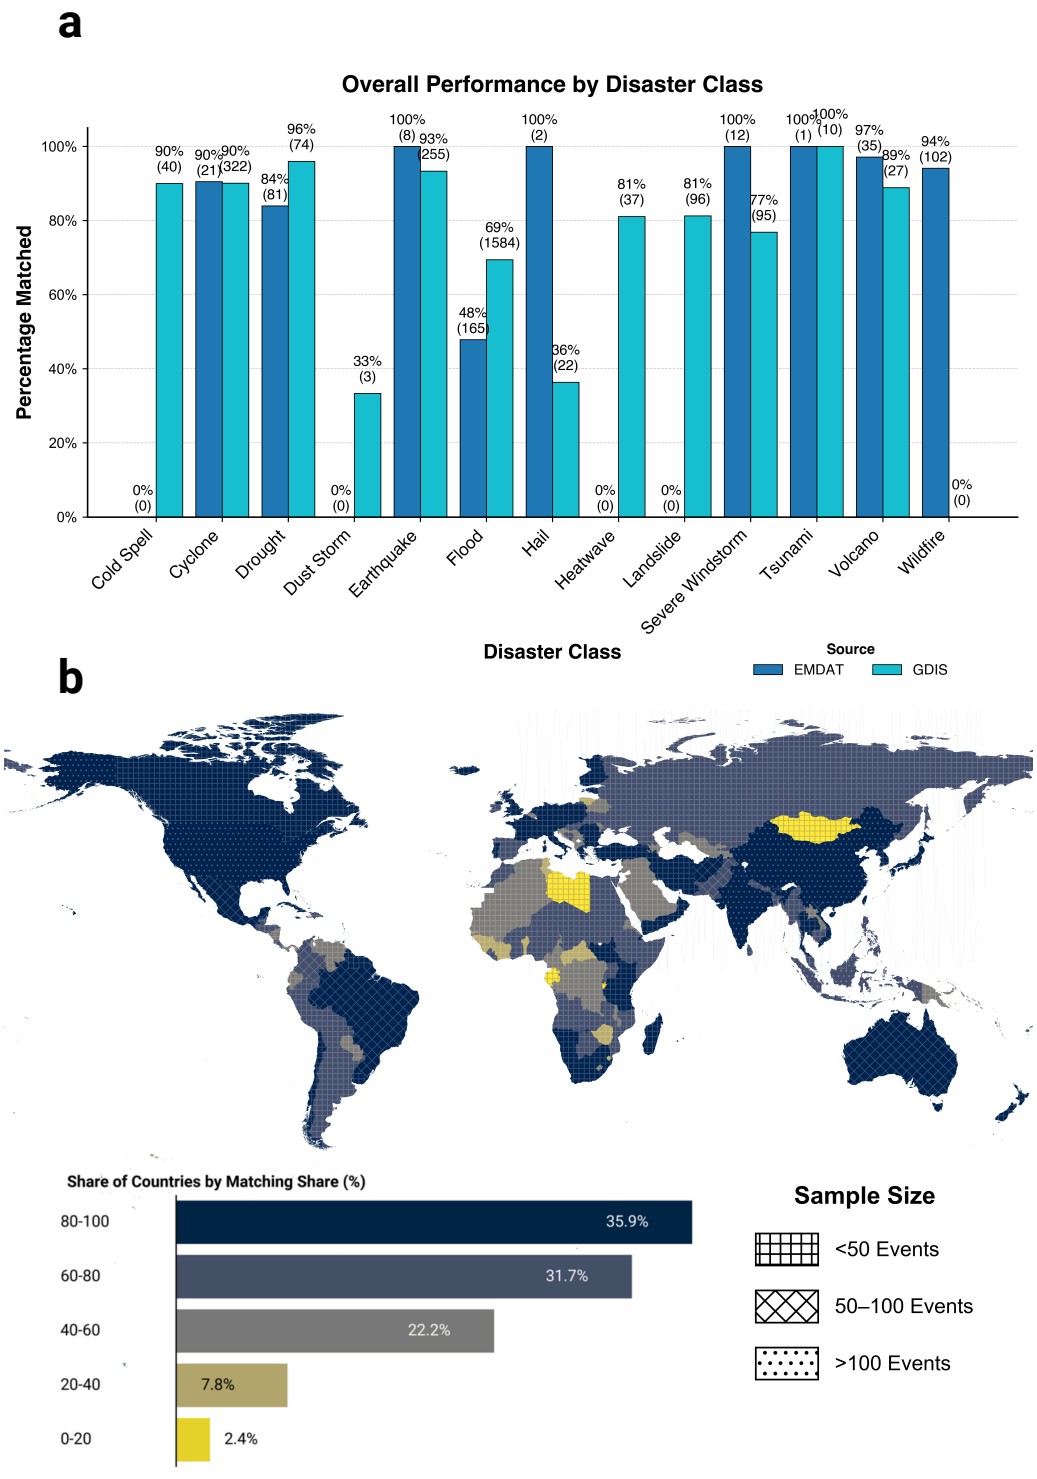

Fig. 2.5. Validation of PEACH data with the EM-DAT and GDIS data sets

Note: Panel (a) shows the matching quality by hazard type and comparative source. A match is defined if for an event at least one cell is treated for the same hazard type and month as long as it falls into any administrative boundaries reported affected in EM-DAT or GDIS. Panel (b) shows the geographic distribution of matching quality along with the number of events by country. The pattern indicates the number of events attempted to be matched and the color of the matching accuracy.

and combine them with the Shapefiles in QGIS, users are advised to, if possible, use R or Python for further analysis to reduce storage space. Information on the scaling of values and missing values can be found in the README file.

The usage is mostly intended for social scientist who want to obtain information on hazard intensities to validate their results. As the data is compiled from peer-reviewed publications from experts in the respective domains, the aim of PEACH is not to serve as a new source of climatological data, but rather to make climatological data more accessible for social scientists. For social scientists many more variables might be of interest. If the data of interest is provided on the country level it can be easily linked to cells using the ISO-3 identifier. This would be the case especially for country-level data from international organisations such as the World Bank, the World Health Organization or the International Monetary Fund. The temporal dimension also allows for the study of structural trends and temporal persistence of effects, which might be relevant for studies on adaptive capacity or climate risk insurance.

Another notable comment refers to the selection process of the included variables. As with any data compilation, decisions were made on why to use this specific data, which might lead to some uncertainty in direct comparison with other sources. This concerns socioeconomic variables, such as population and land use, which are often derived from different sources or employ different methodologies to interpolate the data to a higher resolution. There are ongoing academic debates on the best performing model, which if this is the outcome of interest, it might be noteworthy to see if the used data here is outperformed by more recent models.

Although PEACH has global coverage, it is recommended that it be applied to middle- and low-income countries, where national hazard reporting is more scarce. In western Europe, China, Australia, and the United States more accurate national data is available. In addition, as mentioned in the text, for the two variables: days with thunderstorm environments favorable for hail and days with high dust presence, should be taken with caution, as these two variables are self-calibrated and not derived from peer-reviewed articles. If the primary focus of the study to which PEACH is applied is studying distance-relevant outcomes, one needs to keep in mind that the projection used to create cells is based on the Pseudo-Mercator projection, which leads to distortions in area size as a function of latitude. This means that studies with a primary focus on land area (e.g., through the computation of nightlight per square kilometer) may distort results if the region of interest spans a wide range of latitudes (e.g., in the case of Chile). If equal area and distance within each cell are of key importance

to the study, one might want to re-run the replication code provided on alternative projections such as the Icosahedral Snyder Equal Area Aperture 3 Hexagon (ISEA3H) projection.

The PEACH database is only possible due to amazing research done by the climate, earth and social science community. A list of references to the datasets used is provided within this article and also in the README file of the associated data repository. Please cite when using PEACH also the primary input data to give credit to the authors as stated in this paper to the underlying data sources to give credit to these researchers. The data is provided under the Creative Commons Attribution-NonCommercial (CC BY-NC) license.

2.5 Code availability

The goal of the PEACH platform is to be replicable and accessible to researchers less familiar with Python. The replication files will be uploaded on Zenodo upon acceptance of the submission Reinhardt, 2024b. The code is written in Python 3.9 and run on a server with substantial computational power. Replication requires 600 GB of storage space and at least 16Gb of RAM for most code. While most code is written to be replicable, including the download of the data, some data sets require manual download, which is specified in the README. The code is written in a modular manner and can be easily adapted to other datasets.

B.1 Definition of Cut-Offs

Figure 2.1 of the article shows in Panel (a) the total exposure of each cell to hazards included in PEACH for the period 2004-2015. To build the total sum of events the intensities of each considered hazard need to be transformed into binary dummy variables. The following paragraph will describe the cut-offs used and provide a rationale for each one.

Cold spell events are defined as temperature anomalies that are more than 2.3 standard deviations, so approximately a 99th percentile, below the historic 30-year monthly average. For **heatwaves**, a monthly block maximum from above 33 degrees on the Wet Bulb Global Temperature is taken. This would be equivalent to a black flag within the US Military's Flag system (Patel, Mullen, and Santee, 2013), so a heat level where physical exercise is supposed to be kept to a minimum. **Droughts** are characterized as extreme "compound" droughts, where the definition follows the idea of the intersection of meteorological and hydrological droughts (Wu et al., 2022). A drought month event is defined as one in which the Standardized Precipitation Evapotranspiration Index (SPEI) for both 3-month (meteorological) and 12-month (hydrological) periods is simultaneously below -1.5. **Flood** events are defined as events, where a cell falls within a recorded Dartmouth Flood Observatory flood event (Brakenbridge and Dartmouth Flood Observatory, 2023) with a recorded intensity and at more than 5% of the cell are at high risk of being flooded using the population at risk data mentioned in the paper (Rentschler, Salhab, and Jafino, 2022). A **wildfire** event in a cell is recorded if at least 20% of the cell is burned with a probability of more than 90%. As wildfires are seasonal events that may occur in some regions annually, the deviation of the burned fraction from the expected median burned fraction for the corresponding month is calculated. The median burned fraction is the median observed value for the corresponding month across the entire time period covered. A wildfire event, thus, is finally defined if at least 20% are burned and on average 10% more is burned than the median fraction over the 17-year period (wildfire coverage starts in 2003). A **tropical storm** event is defined if the recorded cell is exposed to winds stronger than 33 meters per second. This corresponds broadly to a Category 1 Hurricane on the Saffir-Simpson Hurricane scale. A **severe wind** event is defined if the wind gusts exceed 30 meters per second on at least one day in the month. Such a threshold, if wind speeds are sustained, would define a violent storm on the Beaufort Wind Scale. However, wind gusts are substantially shorter events. Nonetheless, even if sustained for a short time, winds exceeding this threshold become damaging to infrastructure and pose a risk to

human lives by causing trees to fall. A **hail** event is defined if at least a single day in a month is observed where the probability for hail exceeds 90%. **Dust storm** events are solely based on having at least a single day of what is described in the paper as a "dust environment day". **Landslides** are considered based on the count of events within a cell. There needs to be at least a single event counted. **Earthquakes** are included if the cells have experienced a ground motion with at least scale five on the Modified Mercalli Intensity scale. **Volcano eruptions** are considered as an event if, for the cell, an eruption with at least 1 on the volcano explosivity scale (VEI) is reported. **Tsunami events** are defined as cells having a reported wave height from a tsunami greater than zero. **Extreme air pollution** events are considered if the monthly mean air pollution level exceeds 75 μg/m3 for PM2.5. Based on a recent update by the U.S. Environmental Protection Agency (EPA) this would correspond to an unhealthy air quality index (U.S. Environmental Protection Agency, 2024). As this index is commonly computed on a 24-hour interval, if monthly averages exceed it, one can be certain that these levels are unhealthy for human life.

In this figure, temperature and rainfall anomalies are excluded because of their correlation with flood, heatwave and wildfire occurrences. Using data on population density (Liu, Cao, et al., 2024a) cells below a density of one person per square km in 2019 are excluded.

B.2 Descriptive Economic Analysis

The primary purpose of this article is to explain the methodology and how PEACH was created. The use case of PEACH is versatile and not limited to a specific application. Here, for illustrative purposes, we present two potential use cases, among many, that researchers might be interested in.

National Hotspot Identification

First, the correlation between hazard exposure from 2003 to 2015 and estimated Gross Domestic Product (GDP) at the cell level is illustrated in a descriptive manner. This might help researchers to identify locations where fluctuations in country-level aggregates stem from. Most of the descriptive analysis is based on the data found in PEACH. Yet, PEACH does not include any gridded GDP dataset. To get a proxy of estimated GDP, we rely on the *LitPop* methodology described by Eberenz et al. (2020).

This follows the idea that one can distribute country-level GDP based on a function of population and nighttime light, both of which are variables that we have. The weights by cell are computed here in the following way:

$$Lit^n Pop^m_{gc,2015} = (\text{NTL}_{gc,2015} + 1)^n \times Pop^m_{gc,2015} \tag{2.3}$$

Where *NTL* is the average Nighttime Light value in cell g and country c for the chosen year 2015 and *Pop* the total population in that cell in 2015. The factors n and m are exponents giving the dynamics of the interaction between the two values. The plus one term for nightlight is added to prevent cells with recorded population from being dropped in the event of no available nightlight. Unlike the original version of the paper, we use cell-level aggregates of PEACH here, rather than the pixel-level resolution mentioned in Eberenz et al. (2020). In their article, they test for several combinations of n and m but show that n=1 and m=1 show the highest overall correlation, which is why we also use n=1 and m=1. To create a weight, we compute the GDP in cell g using the following equation:

$$\text{GDP}_{gc,2015} = \text{GDP}_{c,2015} \times \frac{LitPop_{gc,2015}}{\sum_{g \in c} LitPop_{gc,2015}} \tag{2.4}$$

The relative value of the GDP in 2015 of cell g in country c is the product of the total GDP of that country times the relative weight of LitPop over the total sum of LitPop in that country. The GDP data is the value of GDP in current USD from the World Development Indicators of the The World Bank (2024)[2].

Figure B.1 shows the total number of hazard shocks for the most balanced sample from 2003 to 2015 against the relative GDP within each country for the year 2015. All hazards are converted from intensities into binary values using the thresholds explained in subsection B.1. We use the year 2015 for GDP and LitPop, as it is the last year for which PEACH has 100% data coverage. To compute the relative GDP and hazard exposure within each country, we create quantile-based ranks (for three bins) and map the observed values to the tertiles. Cells with null cell GDP are excluded from the map, as this figure focuses on economically relevant areas. Cells with null GDP are areas no population recorded. For an integrated analysis, these areas might still be relevant (e.g. through biodiversity loss), but for this figure, they are excluded. Panel (b) and (c) plot the within-country GDP and hazard exposure ranks, respectively. The figure aims

[2] For Venezuela, the 2014 value is taken given that GDP is not reported for 2015.

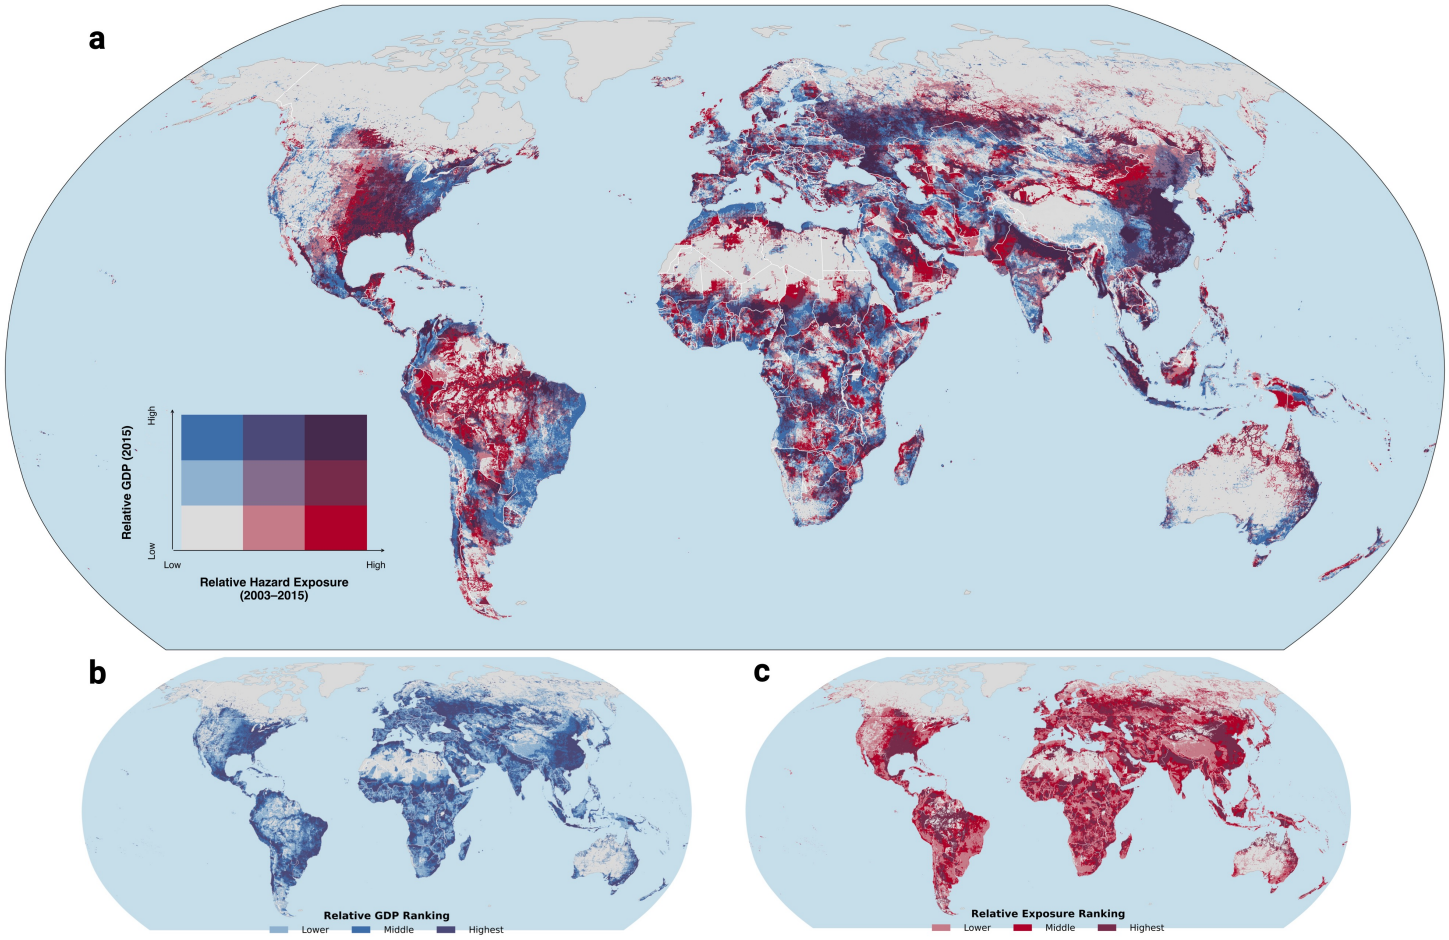

Fig. B.1. National Hazard Hotspots (2003-2015)

Note: Panel (a) shows in a bivariate map centers with relatively high levels of gridded GDP with respect to the rest of the country and relatively high levels of hazard exposure with respect to the rest of the country. Gridded GDP follows the LitPop definition of (Eberenz et al., 2020) and based on values of 2015. Hazard exposure is based on 15 hazard classes between 2004 and 2015. Both are converted into tertile ranks within the country. Gray areas indicate either low GDP and low-risk areas or areas with no population and therefore no recorded GDP. Panel (b) and (c) plot the relative GDP tertile and hazard exposure tertile, respectively.

to illustrate national hazard exposure hotspots. Policy makers might be interested in identifying national areas that see a high number of climate impact drivers, but either high or low GDP. The same analysis can also be performed hazard by hazard or for other outcomes such as critical infrastructure, relative wealth, or, for instance, relative share of agricultural or pasture land.

B.3 Effect of Hazards on Economic Growth

Another application of the data could be to study the interaction of economic growth or nightlight growth, as is often done in economic studies, and cumulative hazard exposure. Quantifying the economic damages of extreme events requires a solid methodological background, as for instance demonstrated in (Burke, Hsiang, and Miguel, 2015; Burke, Zahid, et al., 2024; Hsiang, 2016). Here, for illustrative purposes, we use a simple methodology broadly based on the baseline methodology of Felbermayr, Gröschl, et al. (2022):

$$\Delta \text{NTL}_{gt} = \text{NTL}_{gt-1}\alpha + D_{g,t}\beta + X_{gt}\delta + \delta_t + \delta_g + \epsilon_{gt} \tag{2.5}$$

Where Δ is the first difference of logged nightlight, which measures economic growth using solely nightlight. Note that this differs from the measure used in Figure B.1, where GDP is estimated as the relative weight of country-level GDP by population and nightlight, jointly. Like in most growth equations, the initial pre-shock level of nightlight is included. D is either the total number of shocks of the 14 hazards described above[3] per year or the number of shocks by each individual class, so disaggregated, by year. D additionally includes the lag of each variable. As control, similar to Felbermayr, Gröschl, et al. (2022), we include logged population. Cell-fixed effects absorb region-specific time-invariant factors, while time-fixed effects control for global common shocks. Standard errors are clustered at the cell-year level. This is only a very simplified version, where in contrast to Felbermayr, Gröschl, et al. (2022) neither spatial autocorrelation nor the distinct intensities for hazards is taken into account. The regression is run at the cell level and by geographic regions. The regressions are divided by geographic regions, as the composition of the hazards and economic growth varies widely between these regions. Thus, a comparison pooled over the whole world

[3] Mean and extreme precipitation are excluded as they only have quasi-global coverage. Air pollution is excluded as it is likely endogenous to economic activity and population.

might overlook these differences. Geographic regions follow the United Nations' classification. The excluded territories in the analysis are Greenland and Russia due to their large uninhabited surface areas.

First, the results when D is replaced by the total number of the 14 hazards considered are shown in Figure B.2. One can see that for South East Asia, Central Asia, North America and Oceania, any additional shock always reduces nightlight growth significantly. While for Central America and the Caribbean and East Asia, the effect is positive in the year of treatment, the year after sees larger negative growth rates. For all remaining regions the effect is positive. The results broadly indicate that additional exposure to shocks tends to reduce nightlight growth, except for Europe, the MENA region and South America. This is a descriptive analysis, thus, the effects shown are not causal. Hazard might have strong non-linear effects as found in literature for temperature or rainfall (e.g. see Burke, Hsiang, and Miguel, 2015). Here, the model neither accounts for non-linearity, spatial and temporal autocorrelation nor for the respective intensity of the hazards. However, despite all this, more hazards appear to be harmful to nightlight growth in many parts of the world.

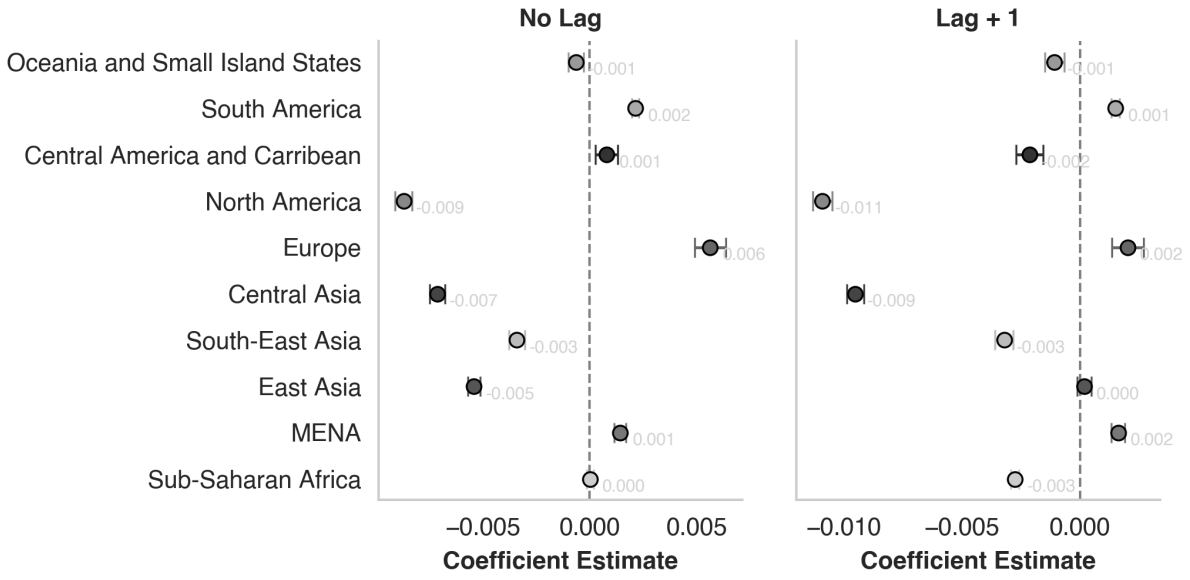

Fig. B.2. Effects of Any Additional Hazard on NTL Growth (2004-2015)

Note: The coefficients show the effects of a single shock more reported in year t (left) and t+1 (right) onto nightlight growth. The model controls for logged population and includes a lagged nightlight control. The individual regressions include time and cell fixed effects, and standard errors are clustered at the cell level. The total shock metric is computed based on a conversion from continuous to binary dummies using the thresholds discussed in subsection B.1.

Finally, the same regression is used to estimate the effect of including all hazard classes individually. Therefore, on a cell level, we maintain the other hazard shocks at their respective regional means. The results are shown in Figure B.3. It becomes directly and visually clear that most hazards, regardless of whether they involve lagged or direct responses, tend to be harmful to nightlight growth. Some hazards always show negative effects on growth independent of the region. This is especially the case for floods, droughts, heat waves, and hail. It also becomes clear that for some hazard-region pairs there are substantially larger effects (for instance, for cyclones and central Asia). This is likely driven by either a small sample size of treated cells or other factors that might drive the growth response. Some regions, such as Europe, seem especially an increase in nightlight in years of additional shocks. Yet, for policy-relevant conclusions, a more in-depth approach should be considered that takes also the interconnections between individual hazards into account and proxies human activity in a more holistic way, accounting for country-specific time-varying factors such as the adaptation capacity, public policies, fiscal stability or social welfare programs.

This descriptive discussion on the economic effects aims to highlight potential use cases of PEACH. One might use it to identify national hotspots that need more attention (relatively low GDP but relatively high hazard exposure). Another potential use case could be to revisit the literature on the connection between economic growth and hazard exposure. The descriptive correlation analysis hinted that the response to shocks varies across geographical regions and across the type of shock considered. Most hazards tend to exhibit a negative correlation with GDP growth; however, a more nuanced and detailed analysis is needed to validate these descriptive findings.

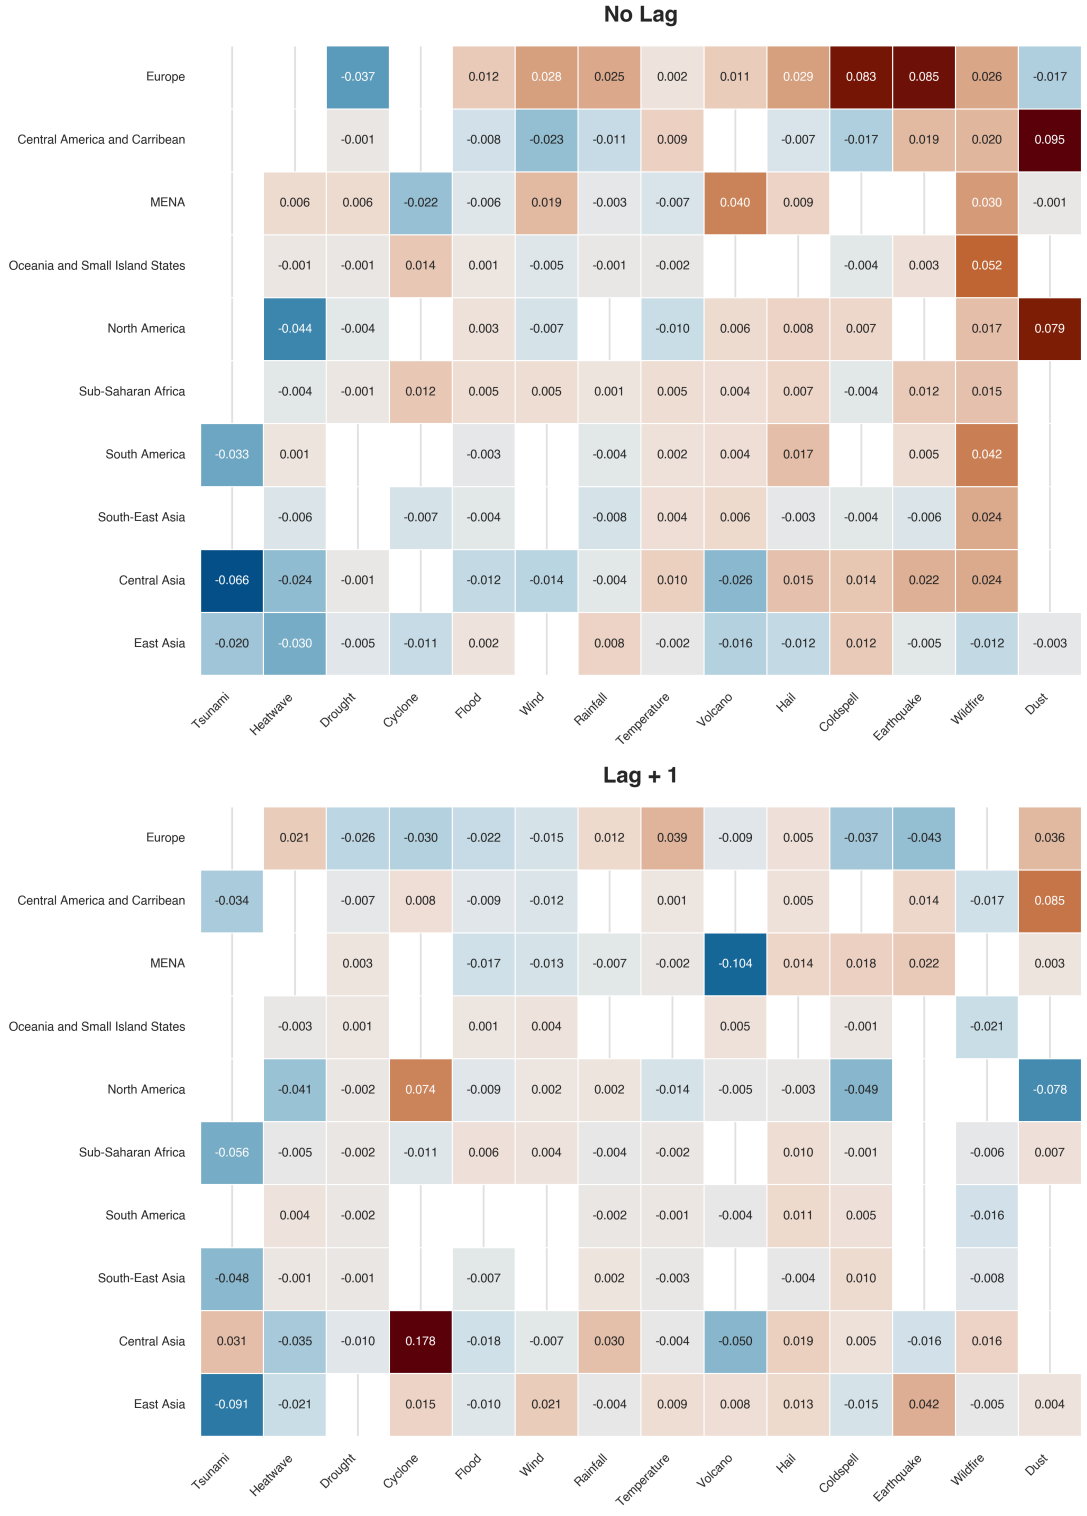

Fig. B.3. Effects of Any Additional Shock by Hazard on NTL Growth (2004-2015)
Note: The coefficients show the impact of an additional hazard shock (x-axis) on nightlight growth in year t (top) and t+1 (bottom). Regressions control for logged population, lagged nightlights, time and cell fixed effects, with standard errors clustered at the cell level. Shocks are converted from continuous to binary dummies using thresholds from subsection B.1. Only coefficients significant at the 5% level are shown; blanks indicate insignificance or lack of data.

CHAPTER 3

Crisis-Driven Mobility: Exploring The Impact of Complex Disasters on Population Changes in Sub-Saharan Africa

This chapter is solely authored.[1]

3.1 Introduction

Extreme natural events are shaping population distributions worldwide, particularly in regions with limited resilience and high agricultural dependence. Regions such as Sub-Saharan Africa (SSA), which fall into this category, see an increase in the frequency and intensity of extreme events (IPCC, 2012). Existing studies have linked climate variability, especially through rainfall and temperature, to agricultural losses and, consequently, to migratory decisions (Hoffmann, Šedová, and Vinke, 2021; Bohra-Mishra, Oppenheimer, and Hsiang, 2014), either at the household level or at aggregate national levels. Far less research is dedicated, however, to understanding how the joint occurrence or accumulation of extreme events affects population change and, thereby, indirectly affects migration trends. This gap is worth investigating, given that such complex disaster environments could have fundamentally different impacts, especially in a context with limited adaptive capacity, such as in parts of rural SSA. Beyond the impact dimension, it is essential to understand how different types of "complex events" - namely single (e.g. one flood), cumulative (e.g. repetition of floods) or "compound" events (e.g. flood and heatwave occurring together), affect net population. This understanding is crucial for anticipating spatial demographic shifts and designing effective adaptation policies.

This article aims to assess the short-term effects of the nine most common hazards in Sub-Saharan Africa on net population change at the local level. Hereby, particular at-

[1] I would like to express my gratitude to Julien Wolfersberger, Rainer Thiele, Stephan Maurer, Lesly Cassin, Kenneth Houngbedji, David Castells-Quintana, Clement Gorin, Sarah Pouey, Sam Marshall and Edouard Pignède as well as the participants from the International Conference on Development Economics, the German Development Conference and French Environmental Economics Conference 2024, the CSAE at Oxford University 2025, as well as seminars and workshops at Paris 1, Columbia university and CIRED for their valuable comments and feedback.

tention is given to studying whether within-year cumulative events and temporally compounding events show fundamentally different responses than aggregate binary treatments. This separation reveals dynamics that aggregate binary treatments conceal, offering a clearer view of how hazards impact population dynamics. We apply an annual panel regression model with unit, time and region-year fixed effects on a sample of 3863 African districts in 44 SSA countries for the time between 2000 and 2019. The main data is data on net population change paired with high-resolution data on monthly hazard exposure for nine hazards (temperature, heatwave, drought, rainfall, flood, hail, wildfire, storm and desert locust). In the first part of this article, we apply a commonly used method for including hazards. Notably, whether or not to have an event. In the second part, we extend this narrow definition and differentiate between the effects of single events, cumulative events (multiple events of the same type in a year), and within-month compound events (overlaps of two different hazards in the same month). Adopting a more standard simple binary definition of extreme events, especially in the case of floods, we document population losses and displacement in affected areas. Similarly, we find an inflow of population following hailstorms in districts that tend to be destinations. The effects of other hazard hazards are inconclusive. In contrast, accounting for all three different treatment regimes, single, cumulative and compound, we find that compound events lead to fundamentally different effects. Although some persistent patterns can be derived for temperature-flood events and flood-storm events, other impacts are context-specific. The first major takeaway is that instead of pooling all events into a single indicator, each hazard combination should be regarded separately, as the impact it has on socioeconomic outcomes may vary.

Moving beyond the classification of extreme events, we examine the context in which hazards occur and propose potential mechanisms behind the documented effects. Most effects are centered around poorer, more vulnerable and remote districts of SSA. The high volatility of the results indicates that the effects are fairly context-specific. As proposed mechanisms, we discuss the effects that the different types of hazards have on agricultural health, crop production, and food insecurity. Here, especially drought and heat-related compound and cumulative events seem to create strong incentives for displacement. The second takeaway is that not only single events matter, but compound and cumulative events are especially significant. This is particularly true for agriculturally dependent rural areas, which tend to show less resilience to these shocks.

The third and final contribution relates to the methodological implications derived from the limitations of this study. Robustness checks reveal the presence of non-

linearity treatment effects and spatial spillovers, which pose significant challenges to causal identification in traditional studies. Rather than treating these findings as technical limitations, we interpret them as evidence of a deeper structural complexity in how hazards affect socioeconomic outcomes. Thus, our model describes dynamic correlations rather than causal treatment effects. This distinction contributes to the literature by emphasizing the need for econometric strategies to jointly account for temporal dynamics, spatial dependencies and latent nonlinearities. Especially in the context of compounding hazards, the article highlights the importance of refining empirical approaches to capture these methodological dimensions of the impact of extreme events.

Our paper aligns with and contributes to two streams of literature: the impact of natural hazards on mobility as an adaptation strategy and the complex effects of hazards. The first stream concerns mobility and migration. The literature on climate migration is well studied and has long been analyzed (Cattaneo, Beine, et al., 2019; Hoffmann, Šedová, and Vinke, 2021). The vast majority of studies focus on panel studies of mostly temperature and rainfall data (for instance, Bertoli et al., 2022; Peri and Sasahara, 2019), with most studies focusing mainly on individual countries and relying on household-level data (Chen, Mueller, et al., 2017; Millock, 2015; Berlemann and Steinhardt, 2017; Gröger and Zylberberg, 2016). Climate-induced migration does not appear to be more prevalent in poorer households and there is great heterogeneity in the type of disaster examined (Kaczan and Orgill-Meyer, 2020). Closely related to this study is the paper by Thalheimer, Otto, and Abele (2021), Norling (2022) and Mueller et al. (2020). The former shows for the East African context with a case study and literature review that for human mobility, especially the context matters, notably the presence of conflicts, social norms, and opportunities for economic improvements. Norling (2022) combines data from the EM-DAT database with data from DHS household surveys in Africa to study the effects of hazards on fertility trends and migration. He finds significant reductions in fertility after droughts and floods. Mueller et al. (2020) use household-level data for four East African countries and study the effects of temperature and rainfall and their interaction on a household's decision to migrate. For their sample, they find little evidence of the rural push of the population in response to temperature or rainfall anomalies. Our study extends previous work by accounting for a broader range of hazard types, encompassing all of sub-Saharan Africa, and explicitly modeling hazard interactions. While we build on Thalheimer, Otto, and Abele (2021) by analyzing co-occurrence, we extend the spatial scope. Like Norling (2022), our outcome implicitly captures fertility through the interpolated net population data, but we actively study

hazard interactions. Similarly to Mueller et al. (2020), we observe that some rural areas exhibit no net population change. In contrast, our study covers a larger scale, covering all of Sub-Saharan Africa, and includes a wider range of hazards.

The other literature to which we add is the nascent literature on cumulative and compounding events. To our knowledge, this is a frontier topic, even in related disciplines such as climate science or impact modeling. If studied, most articles focus on individual hazard classes rather than combining them (see, for instance, Kleemans and Magruder, 2018; Gröger and Zylberberg, 2016; Mahajan and Yang, 2020). However, some studies empirically analyzed the effect of interacting hazard classes, mostly in specific national contexts (e.g. Bohra-Mishra, Oppenheimer, and Hsiang, 2014). Thalheimer, Choquette-Levy, and Garip (2022) adds to this literature by finding that compound droughts increase the probability of migration for the exemplary case of Nepal, Mexico, and Madagascar. Related to this article are the studies by Di Falco et al. (2024) and a current pre-print by Petrova et al. (2025). Di Falco et al. (2024) study the cumulative effects of droughts on migration decisions for five Sub-Saharan African countries. They find that a higher exposure to drought over multiple years increases the likelihood of migration. The pre-print by Petrova et al. (2025) is arguably the closest related to this article. Their study relies on the same data for net population change and studies the effect of crop failures, storms, floods and droughts on net migration. In their baseline, they focus on direct effects of having events, similar to this study, and later study compound events. Although both papers are closely related to this one, given that they rely on cumulative or compound events, there is a key difference. Both studies focus on the "compound" defined as "cumulative over years". This article adds that the short-term response to events might be present even if the events happen within, not across years. Thus, we contribute to the literature by focusing on the interaction of hazards rather than the cumulative effect of events over multiple years, which provides a new perspective.

The paper is structured as follows: First, the data and descriptive statistics are presented in Section 3.2. Then, the methodology of the panel regression model is presented in Section 3.3. In this section, the article also discusses how it defines "compound" events and how they are introduced methodologically. Section 3.4 presents the results of the baseline using aggregate event dummies. Here, also robustness and limitations of the model are assessed. Then, we extend the baseline model in Subsection 3.4.2 and introduce the notion of cumulative and compound events. Section 3.5 investigates the heterogeneity of the sample, followed by the mechanisms discussed in

Section 3.6. The article is concluded in Section 3.7.

3.2 Data

The following paragraphs provide a brief description of the main data sources used in the paper to estimate the effects of natural hazards on population dynamics in SSA.

3.2.1 Dependent Variable

Net Population Change The main variable of interest in this article is net population change. The data is taken from Niva et al. (2023) and shall account for net migration. Although originally presented as net migration data, we refrain from using it as a unique representation of net migration. The data is interpolated from census data combined with birth and death estimates. Net migration is predicted as the difference between the predicted population including fertility and mortality minus the observed population, again based on interpolated population data. These changes in net population might strongly reflect net migration, but not exclusively. Other factors such as census miscounts, vital statistics biases or interpolation biases in the used WorldPop data might also influence the statistics. This might be especially a concern for SSA, where statistical capacity might be weaker (Cameron et al., 2021). Thus, we prefer to call it net population change while acknowledging that large fractions will be explained by migratory movements. The data is provided as a noisy annual 10-km gridded dataset, 3- or 5-year aggregates, or annual district- or province-aggregates. We rely on annual district-level data based on the Database of Global Administrative Area (GADM) Level 2. This allows us to retain sufficient within-annual variation of extreme events. All districts in Sub-Saharan Africa are kept. This represents around 3863 districts during the period from 2000 to 2019. It should be noted that the interpretation of net population change relies on prior knowledge of whether a district is already "losing" or "gaining" population. A decrease in net population would indicate that in "losing" areas more people leave an area, while in "receiving", net-positive areas fewer people arrive. This should be taken into account. Thus, most major regression tables split the sample into "net-receivers" and "net-emitters" as districts that over the 20-year period have seen a consistent outflow or inflow of population.

3.2.2 Independent Variables

The article relies on data from extreme events of nine common hazards in Sub-Saharan Africa, which are deemed important for population changes. The hazard data is taken from Reinhardt (2024a). The original data combines open-accessible data on monthly hazard intensities for 17 hazard classes between 2000 and 2019. Of the 17 hazards, nine are included in this article. The selected hazards are affected by either climate change (drought, temperature, rainfall anomalies, flood, heatwaves, storm) or potentially affect agriculture (hail, locust swarms, wildfires). The data is available at a 10 km resolution for every month. In contrast, the dependent variable is defined at the annual district level. Thus, two aggregation schemes are applied to convert the hazard data to the year-district level. First, we convert the intensities for each 10 km cell into binary dummies using an approach that considers whether a predefined threshold is exceeded. In essence, this takes the following form:

$$D^{(k)}_{gijcrtm} = \begin{cases} 1, & \text{if Intensity}^{(k)}_{gijcrtm} > \tau^{(k)} \\ 0, & \text{otherwise} \end{cases}$$

For each 10 km cell g from Reinhardt (2024a) we compute a binary dummy for each hazard and month. Thus, for each of the nine hazards k we create for each month m, year t, in district i, province j in a country c located in a broader region r (like East Africa) a binary dummy. This dummy is 1 if the intensity recorded for this hazard k in a given month exceeds a specific hazard threshold ($\tau^{(k)}$). The thresholds for each event are specified below. Then, to aggregate the monthly-level binary dummies to an annual level, we count the number of unique sequences. This means that if a drought lasted three consecutive months, it is counted as a single event. If, on the other hand, we have a gap in sequence, e.g. a drought for two months and another one three months later, it will be counted as two sequences. As the duration of events potentially matters, we add robustness tests to assess how our results change using a different temporal aggregation regime. This aggregation gives us an annual count of events per disaster type k at the cell level.

Second, the annual cell-level metric needs to be converted at the district level. Thus, we compute the share of a district affected by each event. To be classified as sufficiently important, a hazard must affect at least 20% of the physical area of a district. This is an arbitrary cut-off point, which may cause secondary problems. For instance, 20% of the area affected in a smaller district portrays a smaller absolute area compared to the

20% area of a large district. Although within a district the relative impact of an event remains the same, the absolute impact across districts can differ. To address this limitation, we use alternative weighting schemes, such as population or cropland affected, in the robustness section. Furthermore, the district-fixed effects, introduced in the next section, reduce the potential bias induced by the cut-off as they indirectly account for the time-invariant area of a district. So for each hazard, a district is considered treated if at least one or more than one event during a year exceeds the intensity threshold τ and covers more than 20% of the district. In the following paragraphs, the individual cutoffs for each hazard (the τ's) are briefly introduced along with the expectation of how each hazard class might influence population change. Original data sources are shown in Table C.1.

Temperature Anomalies The temperature cut-off is based on deviations from monthly historical trends over the last 30 years. To be classified as an "extreme event" a month needs to be in the 99th percentile of historical monthly averages. We expect that temperature anomalies might reduce net population change as it might be harmful to agricultural production, and thus limit possibilities to migrate. However, a rise in global temperature has been shown to be associated with rural-urban migration (Helbling and Meierrieks, 2023; Peri and Sasahara, 2019).

Drought Drought is a slow-onset phenomenon that is potentially harmful to agricultural production. Furthermore, in many parts of SSA, droughts are quite common. To capture upper-tail events, instead of one "extreme" drought, two thresholds need to be exceeded. This follows the methodology of Wu et al. (2022) and defines an event as occurring in a month when the Standardized Precipitation Evapotranspiration Index falls below -1.5 in both a 3-month and a 12-month window simultaneously. This is commonly considered a severe hydrological and agricultural drought environment. Droughts, if they lead to agricultural losses and especially if they accumulate over time, could lead to outmigration of the rural population (e.g., see Thalheimer, Choquette-Levy, and Garip, 2022), thus causing a negative net population change in areas with preexisting population losses.

Heatwave Heatwaves themselves create disamenity and can cause loss of human life, if the local population is not adjusted to the heat (Liu, Qi, et al., 2024; Huang, Masselot, et al., 2023). Although heat waves and temperature anomalies share that both require abnormal temperatures, heat waves mostly strain human health and are influenced not only by temperature but also by factors such as humidity and wind speed. Thus, they are related, but they are separate hazard types. Heat is thus not measured in Celsius

but in the Wet Bulb Global Temperature index. Specifically, as the input data are on a daily level, monthly intensities are the 5-day consecutive block maxima. Given that they are quite common in Western Africa, heatwave events must meet two criteria. A heatwave event must have at least two standard deviations above a 30-year historical mean, but at least an absolute value of 33 degrees. This ensures that the heat event is historically rare and intense enough to significantly strain the human body and pose a tangible risk to health and well-being. We include it for robustness, but it is not expected that heatwave episodes will have tangible impacts on net population change.

Rainfall Extreme rainfall can lead to agricultural gains in arid areas if the soil can absorb the moisture; yet, it can also wash away seeds, resulting in losses and potentially incentivizing reallocation. Rainfall events are defined as 99th percentile events over a 30-year monthly historical mean. This shall capture only the most extreme rainfall anomalies. Rainfall may also induce permanent migration, but to a lesser extent than temperature (see Bohra-Mishra, Oppenheimer, and Hsiang, 2014). Thus, we anticipate a reduction in the local population in the aftermath of extreme rainfall events.

Flood Similarly to heat wave and temperature, flood and rainfall anomalies are related hazards but not directly identical. Rainfall anomalies are, in most cases, the primary cause of floods, especially flash floods; however, other factors, such as dam breaks or storm surges, can also contribute to flooding. In addition, rainfall anomalies may cause flooding not necessarily where the rainfall hits the earth, but rather in downstream riverbeds. Thus, we separate floods and rainfall here, in line with the findings that households respond differently to both types of shocks (Chen, Mueller, et al., 2017). Floods are defined using the Dartmouth Flood Observatory's (DFO) index. The index represents the historical rarity of an event. A flood event is defined as one in which the index reaches or exceeds a value of 1.5. Such floods are rare and typically occur with a probability of 1% to 5% in any given year, which means on average once every 20 to 100 years. This shall minimize areas that are frequently flooded by capturing more extreme events. Floods have been documented to show short displacement effects, but little or no long-term effects (see, e.g. Cattaneo, Beine, et al., 2019 for a discussion). Thus, we expect strong short-term responses to these types of shocks.

Hail Hail commonly destroys crops and therefore can lead to losses in agricultural income. Hail events are defined as months that have at least one day in which hail is very likely. Very likely in our context means a day where the chance of hail exceeds 90%. To our knowledge, there are no studies that connect hail destruction with migration decisions. However, many papers have documented that hazard classes that

destroy livestock and agricultural goods have the strongest impact on a household's decision to migrate (e.g. Blocher, Hoffmann, and Weisz, 2024). Hail may have ambiguous effects, either causing the net population to remain stable if people cannot afford to change locations, or leading to population decreases if they can afford to migrate.

Locusts Locust swarms are the only extended hazard to Reinhardt (2024a) and taken from Food and Agriculture Organization (FAO) of the United Nations (2023). A locust event is defined as having at least one insect that belongs to a locust swarm detected in a district, as it could be indicative of a larger presence of locusts. Locusts are fairly regional in Sub-Saharan Africa, clustered around the boundary zone of the Sahel region. Similarly to hail, we expect that locust swarms do not lead to population changes, despite the swarms destroying harvest, especially when affected individuals are unable to migrate due to income loss.

Wildfire Several parts of SSA face frequent wildfires every year. Wildfire can cause high air pollution and thereby infant mortality (Pullabhotla et al., 2023) or damage agricultural production (Kiely et al., 2021). However, they might also be human-made if used in a slash-and-burn farming practice. Wildfire events are defined as those where, on average, around a fifth of a district is burned. The same concerns apply as previously mentioned regarding the use of a 20% cut-off. In the robustness section, different cut-off points and weighting schemes are assessed. This still leaves a large treated sample. Thus, similar to heatwaves, a secondary condition is added. This states that additionally 10% more of a district needs to be burned than can be expected in an average (median) year over the 20-year period[2]. If wildfires are unanticipated and cause significant crop losses, they can cause the displacement of the local population. Otherwise, no effects are expected.

Storm Storm hazards are grouped here for two distinct hazards: tropical cyclones and extreme wind. Both are short-onset events that can damage the physical infrastructure (Elliott et al., 2019) and cause significant crop losses (Blanc and Strobl, 2016). A storm event is an event where the hourly maximum wind gusts exceed 25 meters per second. This threshold is, for example, found for storm warnings by the European Severe Storms Laboratory (2014). Similarly to floods, storms can have a temporary displacement effect but little or no long-term effects.

[2] As wildfire data is only available from 2001 onward, no historical deviations can be built. Thus, anomalies are built from the median burned area over the 20-year period

3.2.3 Other Variables

In the main part of the article, controls for rural population, cropland share, and conflict are added. Rural population is based on two datasets. First, the district population count is taken from Liu, Cao, et al. (2024a). This data combines common population datasets and fine-tunes them using statistical learning. Next, to identify urban areas, the Global Human Settlement Urban Center Database (GHS-UCD) (Marí Rivero et al., 2024) is added. This data compiles locations and outlines of urban centers in 2024. Urban centers are defined as having dense build-up, a population of more than 50 000 people and exceeding a minimum population density threshold. We assume that most of these cities were already present before 2020. We recompute the population in each urban center and compute the rural population as the difference between the population in the GHS-UCD polygons and the total population in a district.
Cropland shares are computed as area-weighted averages of annual cropland estimates from Winkler et al. (2020). Lastly, the conflict data is taken from the Armed Conflict Location & Event Data (ACLED) database (Raleigh, Kishi, and Linke, 2023) given that it proxies well violent riots and local conflicts, even if no state actors are involved.
In the heterogeneity analysis, we also include data on gridded, time-invariant Human Development Index (HDI) from Sherman et al. (2023) at the district level, calculated as the area-weighted mean. The remoteness of districts is calculated as the average travel time to the market towns with more than 500 000 people from HarvestChoice; International Food Policy Research Institute (IFPRI) (2016). All other datasets, including those used to compute a district-specific climate vulnerability score, are described in their respective sections.

Discussion

The final data set accounts for 3 863 districts in 44 countries and 20 years. Additionally, the data is split into four larger macro-regions (Southern Africa, Eastern Africa, Central Africa and Western Africa) following the United Nations Geographic Regions[3]. The descriptive statistics of the sample are shown in Table 3.1. Columns (2) and (3) divide the sample into areas that, over the 20-year period, lose population (column (2)) or gain population (column (3)). Column (4) tests if these two splits are statistically different. On average, net-emitting regions show higher exposure to tempera-

[3] Accessed 02/07/2025.

Table 3.1
Descriptive Statistics

	(1) Full sample	(2) Net Negative	(3) Net Positive	(4) Difference (3)-(2)
Net Population Change (District)	-1.78	-11.59	12.47	***
	(30.40)	(23.20)	(33.80)	
Hazards				
Temperature: Events	0.54	0.55	0.52	***
	(0.78)	(0.78)	(0.77)	
Drought: Events	0.23	0.23	0.24	
	(0.53)	(0.53)	(0.53)	
Heatwave: Events	0.27	0.25	0.30	***
	(0.66)	(0.63)	(0.70)	
Rainfall: Events	0.31	0.30	0.32	***
	(0.53)	(0.52)	(0.54)	
Flood: Events	0.11	0.11	0.10	***
	(0.31)	(0.32)	(0.30)	
Hail: Events	0.04	0.04	0.03	***
	(0.21)	(0.22)	(0.20)	
Locust: Events	0.03	0.02	0.03	***
	(0.20)	(0.18)	(0.23)	
Wildfire: Events	0.58	0.60	0.54	***
	(0.74)	(0.75)	(0.73)	
Storm: Events	0.03	0.03	0.03	***
	(0.22)	(0.21)	(0.23)	
Controls				
Has Conflict?	0.14	0.13	0.15	***
	(0.35)	(0.34)	(0.36)	
Rural Population (%)	0.88	0.90	0.84	***
	(0.23)	(0.20)	(0.27)	
Land Use: Cropland	0.26	0.31	0.20	***
	(0.31)	(0.33)	(0.28)	
Vulnerability Score (0-100)	70.63	71.59	69.24	***
	(7.50)	(6.64)	(8.41)	
Traveldistance to 500k (hours)	6.58	6.52	6.68	***
	(5.31)	(4.91)	(5.84)	
HDI (mean)	0.46	0.45	0.46	***
	(0.10)	(0.10)	(0.11)	
Observations	77260	45780	31480	77260

Note: p-values: *<0.1,**<0.05,***<0.01. Standard errors are in parentheses. The table shows the descriptive statistics in the overall dataset between the full sample and the sample being split in net-positive and net-negative population data.

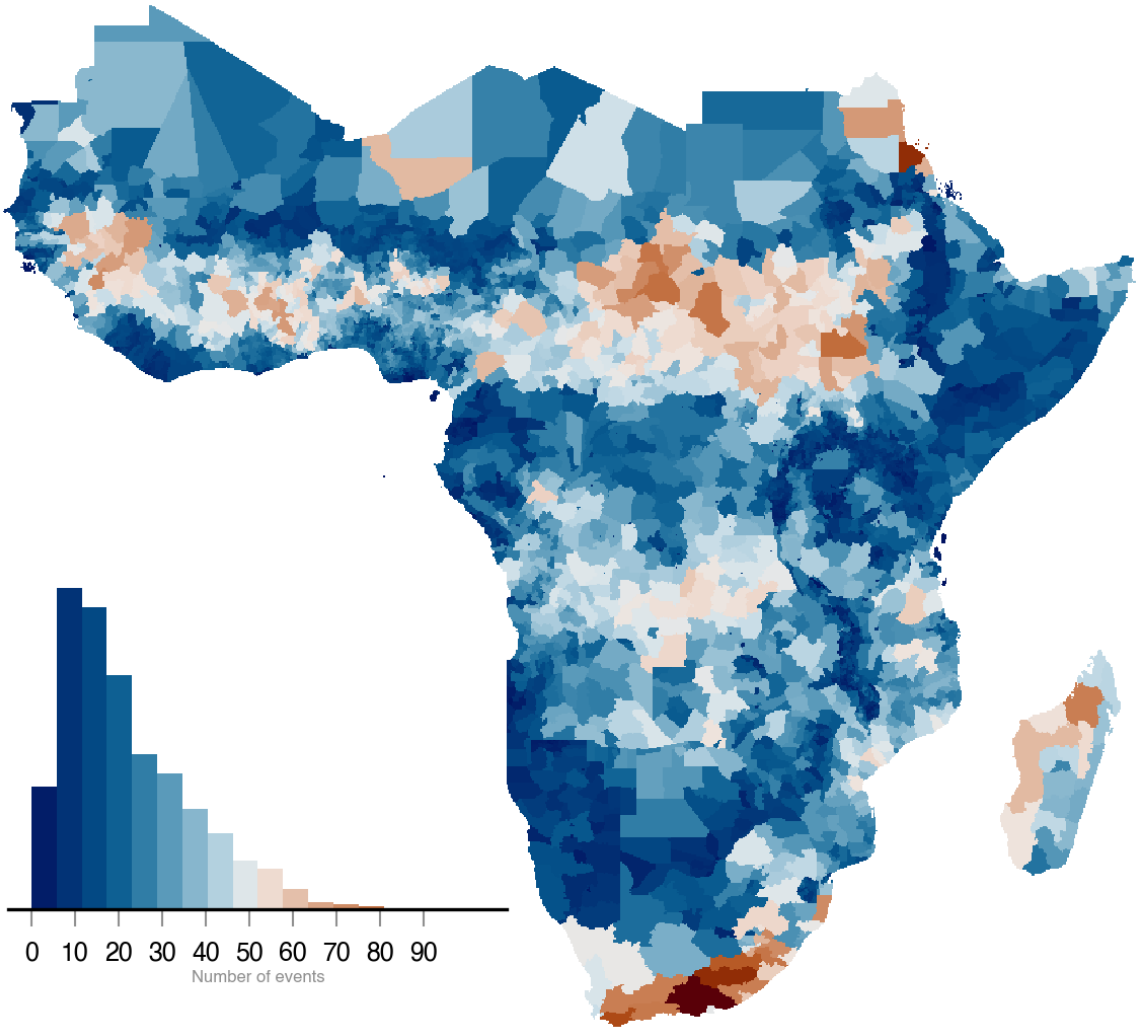

Fig. 3.1. Total Hazard Exposure of Districts in Sub-Saharan Africa

Note: Figure shows the the total number of hazard events of the nine hazards covered in this article at the district-level (based on GADM ADM-2) between 2000-2019. Data is based on, currently unpublished, Reinhardt (2024a). On the lower left, it shows the number of districts and the number of events recorded over the 20-year window. An event is defined as a month-sequence in which each month exceeds pre-specified thresholds. The color scheme is based on Crameri, Shephard, and Heron (2020).

ture, floods, hail events, wildfires, and marginally to storms, while net-receiving areas see more exposure to other events. Net receiving areas also have more conflict, lower rural population, less agricultural land, are less climate vulnerable, but seem to be more remote. That they are more remote is due to some mid-sized cities that have a net positive change but are quite remote in the Sahel region. Furthermore, Figure 3.1 shows that not every district is equally affected by extreme events. The figure plots the total number of events over a 20-year interval and shows that parts of Southern Africa, Sudan, Chad, South Sudan, and Ethiopia are relatively exposed. Most districts during the 20-year period see between 10 and 50 events, in these districts, on average, a negative net population is observed (shown in Figure C.1). To further investigate, Figure C.2 provides more context on the spatial and temporal availability of hazards. In the upper panel, it is evident that temperature anomalies, rainfall anomalies, and drought events are the most widespread throughout the subcontinent. Heatwaves are mainly concentrated in central and western Africa. Floods are more commonly found along the eastern coast. Hail is localized in a few places in Africa. Locust events are only found in the Sahel region. Storms centered around the coastal regions of southern Africa and Madagascar, as well as parts of the Horn of Africa, are likely due to impacts from tropical storms. Most events are available for all years, except for hail, which is only available until 2015. Heat waves and wildfires also miss one year of coverage. Therefore, the data has full hazard coverage for all nine events between 2001 and 2015, and quasi-full coverage thereafter.

The data used in this article is limited in several ways, which has implications for the interpretation of the results. First, the net population data is based on interpolated data from various sources. The statistical capacity of ground-truth data on population, fertility and mortality varies substantially across countries. These differences in capacity are reflected in the quality of the interpolated data. This measurement error, especially in countries with infrequent population censuses, could lead to an underestimation of the effects of hazards on the population, particularly if the interpolated years smooth out short-term variations caused by the hazards. Second, estimates are highly dependent on accurate subnational birth and death statistics reported by national statistical offices. Both fertility and mortality rates might change in the aftermath of extreme events (see e.g. Norling, 2022). If these short-term fluctuations are not captured, our result may miss important short-term demographic responses. Therefore, any significant effects we document are likely rather conservative. If the results document statistically significant effects, despite reduced variation due to interpolation, then the estimates

capture likely lower-bound estimates, as the true magnitude of the coefficients might be larger in reality. To add additional validity to the data, in the robustness section, we aggregate data at the province level, absorbing even more variation. This still broadly confirms the main estimates, providing credibility to the results. However, given these limitations, we refrain from making causal claims.

Second, there might be bias due to the included hazard categories. Most hazard sources are based on reanalysis data, which refers to physical models combined with *in situ* observations, and should adequately mimic global hazards. However, the resolution of the input data varies between 70 km and 300 m. This might lead some hazard classes to be overestimated, especially hail with 70 km resolution. In the 70 km area, not all cells may be equally affected by hailstorms, leading to an overestimation of non-affected cells within the 70 km area. Given the aggregation on a district level this concern is partially alleviated given that the average district size is approximately 7,000 square km substantially larger than the hail resolution. At the same time, hail is only available until 2015, leading to an underestimation of events in the years after 2015. This is a data limitation we cannot avoid, yet we still believe that the inclusion provides additional insights in contrast to previous economic studies. Lastly, the conversion from monthly intensities into yearly event counts may hide substantial intraannual variability and nonlinearity. In the robustness section, it is assessed whether this is indeed the case by using annual maxima of intensities rather than event sequences.

3.3 Methodology

Estimating the effects of treatment in hazardous environments is a challenging task from an empirical perspective. An ideal model needs to accommodate a variation of conditions: First, while depending on the hazard type being seasonal or random in time, a common similarity across hazard classes is the likelihood of reoccurring in the future. Thus, empirical models need to allow for treatment reversal. Secondly, in an optimal methodological setup, one could consider the variability in outcome responses due to individual hazard intensities, coupled with their dynamic interaction with the intensities of other hazard classes. Lastly, following the concept of "constructive destruction" in the literature, disasters can influence the economic growth trajectory of an affected area. An ideal model takes this temporal dynamic into account.

3.3.1 Baseline Model

This article relies on time-dynamic treatment models in a panel with three sets of fixed effects to establish a connection between hazard exposure and net population change. Although far from perfect, this type of model allows, under fairly strict assumptions, for treatment to recur multiple times and for interactions between different hazard classes to be accommodated. The baseline model is build on the following equation:

$$Y_{ijcrt} = \beta_0 + \beta_1 D^{(k)}_{ijcr,t-1} + \sum_{\substack{l=1 \\ l \neq k}}^{L} \beta_l T^{(l)}_{ijcr,t-1} + \gamma X_{ijcr,t-2} + \lambda_i + \lambda_t + \lambda_{r \times t} + \epsilon_{jcrt} \tag{3.1}$$

The outcome Y measures the net population change per 1000 people in district i, province j, country c, geographic continental regions r (West, East, South and Central Africa) and year t. The $D^{(k)}_{ijcr,t-1}$ is a binary dummy that is one if at least one event of type k took place in $t-1$. The model is rerun for every specific hazard k. Treatment is defined as having an event sequence exceeding pre-specified thresholds in at least 20% of a district. In contrast to previous studies, the model also controls for "background stressors". These are entered as T variables and refer to other hazards that exclude k ($l \neq k$) that could also have occurred. They shall reduce the risk of omitted variable bias and serve purely as controls. These variables are added as additive cumulative counts of all other hazards that are not k. Thus, each cumulative count for the L hazards has its own coefficient. The vector X represents the control variables. In particular, it accounts for the share of the rural population, the share of cropland, and a dummy that has had a violent conflict. Rural population aims to proxy the "rurality" of a district. Migration and population change might be an important migration strategy for the rural population (Beine and Jeusette, 2021; Hoffmann, Šedová, and Vinke, 2021). Similarly, a higher agricultural dependence, measured by the share of cropland, could influence a migratory response through agricultural income (Blocher, Hoffmann, and Weisz, 2024). Conflict, especially in connection with extreme events, might influence migratory decisions, and therefore should be taken into account (Thalheimer, Otto, and Abele, 2021). Thus, these three controls are included to account for the influence of these categories. However, following Dell, Jones, and Olken (2014) and Hoffmann, Šedová, and Vinke (2021) these controls might represent bad controls, as rural population and conflict are likely strongly correlated with natural hazards and obfuscate somewhat the direct effects of hazards on migration. To minimize the risk that they serve as mediators for

the treatment effect, they are lagged by two periods. The term λ_i accounts for district-fixed effects. They control for time-invariant district-specific characteristics such as geographic characteristics (mean elevation or proximity to borders) as well as cultural and historical norms that might be important for fertility and population trends. Time fixed-effects (λ_t) control for joint common temporal shocks such as global commodity price shocks. Inclusion of these fixed effects are common in the literature (see e.g. Dell, Jones, and Olken, 2014). Additionally, we control for sub-continental region times year fixed effects ($\lambda_{r\times t}$). These sub-continental region (e.g. West Africa) are interacted with annual dummies to control for macro-climatic variation. Many large-scale climate phenomena, such as the well-known El Niño, often show similar characteristics in broader continental regions. For example, it might cause widespread drought conditions in Southern African countries, but not equally across all of them. If not controlled for, these factors remain in the error term and potentially influence the occurrence of extreme events, thus serving as omitted variables. Standard errors are clustered at the province level to allow for the within-province correlation of the error term. Provinces are defined based on the Database of Global Administrative Areas (GADM) Level 1.

Identification

The main coefficient of interest is β_1. It measures the short-run, homogeneous marginal effect of being exposed to a certain type of extreme event controlling for contemporaneous cumulative exposure to most other hazard events. Identification in this model originates from the random treatment timing and relies particularly on the strict exogeneity assumption. The method assumes that each treatment, conditional on the controls, background stressors and fixed-effects is orthogonal to the standard errors over all periods. Following Imai and Kim (2019) this implies that there are no excluded time-varying confounders, there is no feedback from past outcomes to current treatment nor carry-over effects meaning that past treatment affect current outcome. Lastly, past outcomes are independent of current outcomes. Although extreme events are plausibly exogenous in their occurrence, violations of strict exogeneity may still arise if events are anticipated (e.g., via early warning systems) or have persistent effects that extend beyond the treatment window. The next section discusses threats to identification, including those that challenge the strict exogeneity assumption.

The first and most discussed issue is the concern of model misspecification, especially through omitted variable bias. If unobservables remain in the error term and have an effect on net population change and treatment or controls, the effects are biased. Following Hogan and Schlenker (2024) this is of particularly concern if treatment de-

scribes short-run weather events if variables that belong to the same distribution of the weather event are not included. An example would be agricultural studies that use temperature as a treatment, but omit the consideration that soil moisture could also have an impact on the outcome. We believe that this is of a lesser concern in this study, given that each model controls for the presence of other hazards, potentially drawn from the same district-specific distribution, and by the inclusion of the three sets of fixed effects. Furthermore, given that the events are aggregated at an annual level, much of the variation is lost. However, as in many macro-level migration studies, a small concern remains regarding omitted variable bias. Furthermore, our model might be misspecified due to not capturing non-linear effects. In the current model, homogeneous treatment effects are estimated. However, hazards might have important non-linear effects, which are left unaccounted for in Equation 3.1 and remain in the error term. In the robustness section of the baseline, we test in a simple manner whether, when using duration instead of binary treatments, duration enters nonlinearly for each hazard.

A second concern is the violation of the Stable Unit Treatment Value Assumption (SUTVA). It assumes that being treated is consistent across units and there are no spillovers, peer effects or other unobserved heterogeneity. We assume that anticipation plays only a minor role, given that early warning systems are not universally available in Sub-Saharan Africa, especially in rural, underserved districts. However, temporary displacement of populations and long-term decisions to move after hazards might cause rural-urban migration (see, e.g. Adger et al., 2020). This would cause spatial spillovers. To assess this, we run a Spatial Durbin Error Model (SDEM) in the robustness section and find that, unsurprisingly, spatial spillovers are indeed found for some hazards. Thus, while Equation 3.1 assumes that SUTVA holds, in reality, there is a risk that part of the captured effects are driven by spillovers from nearby districts.

Third is the concern on selection into treatment. Natural hazards may be random in time, but not necessarily random in space. As Figure C.2 visualizes, especially the storm and locust events are clustered among certain parts of Africa. Using the full sample of untreated units leads to a comparison of units that could vary substantially in their respective risk of exposure. Thus, we add a robustness check to see if the results are comparable, accounting only for units that have at least one treatment over the 20-year window.

Fourth, there is the concern of temporal heterogeneity. As Equation 3.1 only captures a one-year lag, it omits other dynamic treatment dimensions as it assumes treatments

take a year to fully influence net population change. However, dynamic effects might have shorter or longer effects than just a single year. If excluded, these dynamic treatment variables will remain in the error term, potentially biasing the results. Applying one of the more recent event study estimators, among others, those by Callaway and Sant'Anna (2021), Borusyak, Jaravel, and Spiess (2024), and Wooldridge (2021) may be a solution to the problem. While allowing for the capture of dynamic effects and the establishment of valid control groups, despite their use in the literature, we argue that they do not fully apply to our setting. This is notably because, in most cases, units that are treated once will be treated multiple times over the 20-year horizon. Baker et al. (2025) discuss that predominantly three models allow for such treatment reversal, notably the work by De Chaisemartin and d'Haultfoeuille (2024), Imai, Kim, and Wang (2023), and Liu, Wang, and Xu (2024). The estimators rely on different assumptions (e.g. limited carryover assumption as Imai, Kim, and Wang (2023)) to tackle the issue that there is not a unique treatment but rather a unique sequence of treatments (e.g. 0,1,0,1,1). To our knowledge, the least restrictive estimator allowing for reversal is the estimator by De Chaisemartin and d'Haultfoeuille (2024), which will be used here for robustness. While fairly flexible, this estimator still relies on a generalized version of the parallel trend assumption, specifically that if two districts are treated at the same time, their expected path on the outcome would be the same in the absence of the treatment (in our case conditioned on the fixed effects and controls). This assumption likely fails in reality, given that net population change and net migration follow a multitude of influencing factors such as social networks, cultural norms or crop diversification, each of which could cause deviations from parallel trends. In addition, every event might lead to different growth trajectories causing diverging trends over time. This challenges causal identification, particularly in disaster studies as the number of clean "pre"-trends shrinks. However, we include this robustness test to evaluate whether temporal dynamic effects drive the results in our baseline model. Furthermore, although strong, we assume that the conditional parallel trends assumption holds in the baseline specification. That is, conditional on district, year, and region-year fixed effects, together with included controls, treated and untreated districts would have evolved similarly in the absence of an event of type k. Lastly, as the baseline model does not rely on a causal estimator, we do not claim causality.

Finally, the estimated model might suffer from covariate bias, meaning that treatment effects are explained by covariates included. This is especially a concern if the treatments themselves are strongly correlated. Figure C.3 shows that the correlation exists,

but is with a maximum of 0.3, which is never extremely large[4].

In general, identifying the causal effect of hazards on outcomes is a difficult task. Not only does each hazard influence population trajectories through a variety of factors, but the recurrence of hazards, their spatial spillovers, and the potential nonlinearities of impacts make it challenging to truly claim causality. Along with the interpolated dependent variable, this article does not claim causal identification. Instead, it provides insights into the dynamic correlation of hazard events on net population change.

3.3.2 Extended Model

The estimated Equation 3.1 substantially oversimplifies the real-world environment. The coefficient D is shown as binary in having at least one event of class k. However, in practice, having one or having five events might have fundamentally different impacts. Alternatively, effects might also become stronger if multiple hazard types co-occur. Mostly due to the progress of global warming, climate stressors occur more frequently simultaneously or in rapid succession to each other (Seneviratne, Zhang, et al., 2021). In this extension to the baseline model, we will separate these three treatment regimes. This follows broadly the notion of Zscheischler, Martius, et al. (2020) that discusses so-called compound events. They discuss different forms including temporal compounding events (the succession of various events in the same region) or spatially compounding events (multiple events happening at the same time but at different places). These events can be of the same type or of multiple types (so-called "multivariate"). We focus primarily on the notion of temporally compounding events here. The three treatment regimes are therefore: first, *single events* are hazards where only a single time within a year window a sequence has been detected, where thresholds are exceeded. The creation is identical to the one used in the previous part. An example is a storm that is recorded in a given year. Second, *cumulative* events are hazards of the same type when more than one sequence is recorded in a given year, in all thresholds have been exceeded. Again, creation is conducted as before. An example would be a district that sees two floods in a given year. We do not distinguish between the timing between each event, which might also be important. However, to be considered, there must be at least one month between each sequence where the threshold is not exceeded. Third, *compound* events are defined as the joint occurrence

[4] The figure does not show the significance of correlation given that mechanically a large number of observations will almost always cause Pearson correlation to turn significant.

of two different hazard types within a one-month window, where both events exceed the predefined thresholds. More details are provided in subsection C.3. An example is the common co-occurrence of a flood and storm during the landfall of a tropical storm. To avoid double counting, single events and cumulative events are mutually exclusive, as well as, if a compound event is recorded, this sequence is subtracted from the other two events. This means that if a year experiences both a flood and a storm, but they overlap, neither the flood nor the storm is recorded individually; only the compound flood-storm interaction is included. An overview on the hazards is shown in Table C.2. Looking at the fraction of events by year for any of the three types in Figure C.4, such "compound" events account in most years for more than half of all events motivating the study of their impact more in-depth.

To include all these coefficients, a full extended model is specified as:

$$Y_{ijcrt} = \beta_0 + \underbrace{\sum_{k=1}^{K_1} \beta_k^{(1)} D_{ijcr,t-1}^{(k,1)}}_{\text{Single events}} + \underbrace{\sum_{k=K_1+1}^{K} \beta_k^{(>1)} D_{ijcr,t-1}^{(k,>1)}}_{\text{Cumulative events}} \tag{3.2}$$

$$+ \underbrace{\sum_{k=1}^{K} \sum_{\substack{l=1 \\ l \neq k}}^{L} \beta_{kl}^{(\text{int})} D_{ijcr,t-1}^{(k,l)}}_{\text{Multivariate interactions}} + \gamma X_{ijcr,t-2} + \lambda_i + \lambda_t + \lambda_{r \times t} + \epsilon_{jcrt} \tag{3.3}$$

The updated model retains the same fixed effects and controls as the baseline model. However, instead of having only a single treatment dummy, all coefficients are added at once. $D^{k,1}$ is a dummy for each of the nine types of hazards and equals one if only a single hazard has been recorded in a given year. $D^{k,>1}$ is a dummy equal to one if more than one event of hazard type k has been recorded in year $t-1$. Finally, $D^{k,l}$ provides all unique combinations of pairs of two different hazard types and equals 1 if in a given year two hazards co-occurred at the same month.

The adoption of this model goes beyond simply showing that single events, cumulative events and "compound" events are different. We expect that single events ($\beta_k^{(1)}$) have weak effects on net population change, especially through migration[5]. Some events, especially destructive ones might force people to migrate, similar to those that destroy sources of income. Cumulative events ($\beta_k^{>1}$) might amplify the direct effects.

[5] Given that net population change is $P_{t+1} = P_t + B_t - D_t + M_t$ where P is population, B is birth rate, D is death rate and M is net migration and B and D evolve usually rather slowly, large parts of net population changes can be attributed to migration.

If, however, areas with repeated exposure are more adapted to these shocks, cumulative events might also reduce net population change. Households might be able to absorb a single event, but repeated events might exceed the maximum welfare losses that households might be able to absorb. A discussion on such capacity can be found in Hallegatte (2014). Finally, compound events ($\beta_{kl}^{(int)}$) might have fundamentally different effects. The coefficients might be negative if two co-occurring events amplify the damages. For example, a storm and a flood, commonly found in the landfall of tropical cyclones, may amplify the damage caused by the wind alone. In contrast, some combinations might offset each other's effects on the population. For example, a rainfall shock in the Sahel zone following a locust outbreak could help offset some of the negative effects on crops, thus limiting the effects to if the locust swarm had been present alone. This complexity that events might amplify through cumulation but be offset in case of interactions motivates us to study these three separate treatment regimes.

The same identification concerns discussed above apply. Additionally, there are two further noteworthy concerns. First, as shown in Table C.2, some combinations are extremely rare. Thus, in the final regression model, events with less than 30 treated units are excluded. These hazard pairs are likely statistically underpowered, causing wrong test statistics, but they increase the likelihood of overfitting the model. Second, having 48 coefficients in the model risks overfitting and causing imperfect multicollinearity, which means strong correlation between the individual coefficients that inflates standard errors and obfuscates treatment effects. Thus, as robustness, we apply the Post-Double Selection Lasso regularization of Belloni, Chernozhukov, and Hansen (2014) and implemented by Ahrens, Hansen, and Schaffer (2024) on each coefficient iteratively. This will retain for each treatment dummy only those variables that are jointly important for both the net population change and the treatment dummy. The procedure shall ensure a valid variable sample for inference that minimizes the risk of overfitting.

3.4 Results

Having introduced the baseline and the extended model in the following pages, we will present the baseline results, interpret them, and test for their robustness. Next, as described above, we will split the binary treatment dummy into three distinct parts and present the results. Finally, the extended full model tests for heterogeneity and tests four proposed mechanisms.

3.4.1 The Effects of Being Exposed on Net Population Change

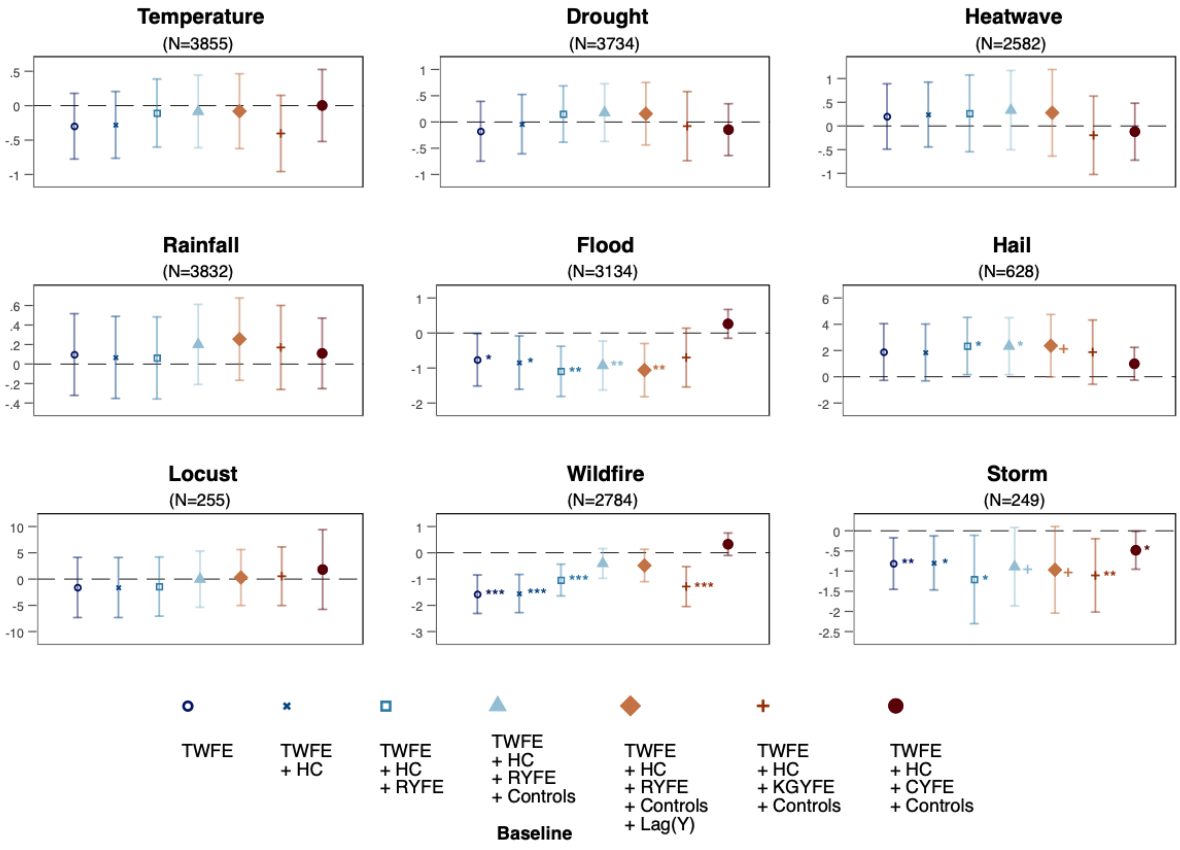

Fig. 3.2. Baseline Estimates by Hazard Type

Note: *** 0.01, ** 0.05, * 0.1, + 0.15 pval. The figure shows by each of the nine hazard types the results of Equation 3.1. It moves from less restrictive models to more restrictive ones. TWFE: two-way fixed effect with district and year-FE. HC: Hazard controls, which control for cumulative background stressors. RYFE: Continental Region-year fixed effect. Controls: control for 2-year lagged rural population, cropland share, and conflict dummy. Lag(Y) lags net population change by one period. KGYFE: replaces continental region-year FE with agriclimatic Koppen-Geiger times year FE. Data is taken from Beck et al. (2023). CYFE: country-year FE models. The N indicates the number of districts that ever have this event. All models cluster standard errors at the province level.

The baseline model, along with different model specifications, is shown in Figure 3.2. The models start with simple two-way fixed effects (the hollow circle on the left) and gradually add background controls (the small x), region-year fixed effects (the squares) and finally the other two-year lagged controls for rural population share, cropland

share, and conflict presence (the triangle). This is the baseline model described in Equation 3.1. Of all nine hazards, only flood, storm, hail and wildfire show any significance. Floods reduce the net population change significantly by around 1 per 1 000. Hail becomes only significant after controlling for region-year FE. Hail events tend to increase net population change by around two persons per 1000. Wildfire has large and significantly negative coefficients until controls are added. This might suggest that areas prone to wildfires are also likely to experience conflict or significant fluctuations in the rural population. Similarly, the storm remains significant at a 10% level until controls are added, reducing the coefficients to only significance at a 15% level, below standard norms. Next, we add additional specifications. First, we add a lagged dependent variable in the regression model to turn the regression model into a dynamic model[6] shown as a rhombus. The results are comparable to the baseline. Next, instead of continental region-year fixed effects, agriclimatic zone-year fixed effects are added, indicated by the plus symbol. These zones, based on the Koppen-Geiger classification, shall account for large-scale agricultural suitability. Interacted with time, it absorbs variation specific to these agriclimatic zones. It is used as a robustness to geographic region-year fixed effects assuming that instead of geographic regions, large-scale atmospheric conditions might influence agricultural zones rather than geographic ones. Flood becomes insignificant, as does hail. In contrast, wildfire and storm turn significant again. This means that likely part of the variation originally captured by floods and hail is due to large-scale geographic trends. Wildfires and storms can affect larger areas within geographic regions, rendering coefficients insignificant once their variation is absorbed by the fixed effects. Agricultural zone times year fixed in contrast can vary largely within each geographic region, leaving more variation to make the coefficients significant. Lastly, we also add country-year fixed effects, which is indicated by the round circle. Although this seems a natural way to control for country-specific variation, a problem is that some hazards have a broad spatial scale. While country-year fixed effects work well for countries that stretch extensively (such as Nigeria, South Africa or Mali), for small countries (such as Burundi, Rwanda, Lesotho) most of the annual variation is absorbed leaving none to identify treatment effects. Thus, we exclude country-year fixed effects and continue with region-year fixed effects.

Net population change is complicated to interpret as it depends on whether areas gain

[6] We are aware of the risk of the so-called "Nickell bias" that might arise from adding a lagged dependent variable. However, we do not believe that this will fundamentally change the results given that the time dimension is fairly long

or lose population. For instance, a decline in net population could mean two possible channels: in areas that see on average more people moving in, it indicates a reduction in the number of people moving in. In contrast, areas that are losing local population would mean more people leaving the area. Thus, to better understand whether results are indicative of in-migration or out-migration, the sample is split into districts that are net receivers over 20 years (with positive net population change) and net emitters (with negative net population change). The results are shown in Figure C.5. Splitting the sample reveals that, at a 10% level, only flood and hail remain significant. Floods lead to negative net population change, especially in net-emitting regions. Thus, floods appear to lead to out-migration, potentially through temporary displacement. Hail leads to an increase in net population change, indicative of an increase in in-migration. A conjectured hypothesis here is that if hail destroys crops, it might incentivize people to relocate to economic centers that absorb the farmers who lost their income due to hail. In general, the results point to a surprising picture in addition to floods and hail, and none of the expected hazard events show substantial effects. Later, we will highlight why this might be the case.

Robustness of the Baseline Model

In the following, several robustness checks are conducted for the baseline model. The results are shown in the Appendix.

Functional Misspecification First, we assess whether the functional form used is correct or misspecified. At the same time, it is tested whether using sequences instead of monthly hazard counts changes the baseline result. Table Table C.5 shows the use of monthly counts rather than sequences as treatment variables. Flood and hail show a similar direction as in the baseline, but also show nonlinear effects, where more prolonged floods turn coefficients more positive, while the opposite is true for hail. Prolonged exposure to temperature appears to have a linear effect. The turning points to see declining effects is only found for around a quarter (floods) and less than 10% (hail) of all events. This means that the baseline holds for the majority of events in terms of functional form. However, for most upper-tail events, some of the effects found might be reversed.

Violation of the SUTVA Assumption The model could be biased if anticipation is present or spatial spillovers exist. Although anticipation is difficult to assess in the current setup, we assess bias from spatial spillovers. Spatial spillovers from natural

hazards have been documented several times in the literature (e.g. Conigliani, Costantini, and Finardi, 2022; Felbermayr, Gröschl, et al., 2022). To test for this violation, we run a Spatial Durbin Error Model following Anselin (1988). This enables us to distinguish between direct effects and spatial spillover effects. The results are shown in Table C.6. There are documented spillovers, especially for temperature anomalies and droughts. Where some of the effects are offset, if neighboring districts are also affected. Floods seem to matter mainly locally. Locust swarms, given the high spatial clustering around a few parts of Africa, become strongly significant, where shocks cause large increases in net population change. Lastly, accounting for spatial spillovers hail becomes insignificant, highlighting that some hidden spatial dynamics might drive the results. Overall, this suggests that SUTVA may be violated due to the presence of spatial spillovers. Thus, our baseline estimates partially capture the spillover effects from neighboring districts, especially for temperature anomalies and droughts. Thus, the proposed findings must be interpreted with care and not as causal, but rather as a correlation.

Selection into Treatment In the baseline model, the chosen comparison group also includes districts that are neither at risk nor ever exposed to the specific hazard under consideration. To see if the results change, we condition the regression on having had at least one event of the hazard over the 20-year period. The results are shown in Figure C.9. The results are comparable and become slightly more significant, especially for hail.

Micro-sample Validation To validate the results, micro-level data on migration decisions of households are added. To this end, the World Bank's LSMS-ISA data for Nigeria is utilized, capturing three waves of household surveys. Then, using a logit model with controls similar to Di Falco et al. (2024), the effects of having had a hazard event on the likelihood of having a migrant are investigated. The results are shown in Figure C.11. In particular, floods and droughts significantly reduce the likelihood of migration. Given that this translates to reduced out-migration from a district, this means a positive net population change. Hail has little to no effect because only a few parts of Nigeria are exposed to these types of events. This test highlights that especially drought and flood influence migration and thus net population change. Hail might still have effects on the whole continent sample, but for the specific case of Nigeria this does not seem to be the case. In addition, it is discussed what happens if the shocks accumulate over time. Instead of "within-year" cumulation as in the rest of the paper, here, "across-year" cumulation is studied. This means what happens if a certain hazard

appears multiple years in a row. It is found that, especially for temperature anomalies, these types of cumulative multi-year events have rather large effects. Unlike the articles by Petrova et al. (2025) and Di Falco et al. (2024), this type of cumulative event is not at the center of this article. However, we acknowledge that the effects might amplify once the repeated exposure is taken into account over multiple years.

Dynamic Treatment Effects The baseline model assumes that treatment takes a year to impact net population change. However, the impacts are likely more direct or take more time to manifest. To assess this, a dynamic event-study estimator by De Chaisemartin and d'Haultfoeuille (2024) is used on each of the nine treatment dummies. This estimator is one of four current difference-in-difference estimators that allow treatment to be turned on and off over multiple years. The results, shown in Figure C.10, reveal two main findings: first, anticipation of treatment appears to be limited to temperature anomalies, where clear pre-trends are observed. Second, our baseline may miss dynamic carry-over effects for wildfires and storms that start one year after the event and continue onward. Thus, a minor part of our documented coefficients for these two events might be driven by carryover effects from previous treatments. For floods, only the year after the event is significant, confirming the baseline results.

Other Robustness Checks Additionally, it is assessed if large outlier districts drive the main results by using a 1 and 5% winsoring of the outcome variable in Figure C.13. This appears not to be the case. Next, it is assessed whether the secondary condition that at least 20% of a district needs to be affected by changes results in Figure C.14. It appears that the significance of hail is marginally dependent on the threshold picked, but broadly holds. Lastly, Figure C.15 aggregates the results at the provincial level to see if results hold on a broader geographic scale. Hail becomes insignificant, but wildfires, flood and droughts show significant changes in net population, especially in net-receiving (droughts) and net-emitting (wildfires) provinces. This hints that these hazard have likely broad-scale impact dimensions that affect population beyond the district-level. Hail, in contrast, might have more local effects.

The robustness checks aim to stress test the validity of the baseline model. They reveal evidence of nonlinearity for a small subset of events and hidden spatial spillovers for temperature and droughts that are not accounted for in the baseline. They thereby emphasize that the chosen baseline model simplifies reality. This reconfirms the initial assessment made that the findings presented shall serve as suggestive patterns rather than causal estimates. However, these robustness tests reveal two important takeaways: first, the spatial interdependence reveals a migration and displacement story.

Changes in net population in one district impacted by shocks in another indicate the presence of migration in response to shocks, being part of a migration-as-a-response story to shocks. At the same time, the nonlinearities reveal in parts that the response to exposure depends on its intensity, meaning that different levels of asset or crop destruction might lead to different population change responses. Some of the violations found actually hint strongly at hidden mechanisms. Some of which will be tested later. Second, these tests stress that future models need simultaneously to account for nonlinearities, spatial spillovers, as well as the dynamic treatment dimensions. To our knowledge, no current model allows us to account for all three simultaneously. However, they also confirm that natural hazards do cause changes in net population, and the presented effects are averaged effects. In the following paragraphs, we will elaborate on another aspect of the model, specifically the treatment itself.

3.4.2 Extending the Baseline: Introducing Three Treatment Regimes

Up until now the models presented are build on Equation 3.1. They account for a treatment of interest, where an event is defined as having occurred at least once in a year. This oversimplifies reality. It may matter whether you experience a single flood, multiple floods in a year, or a combination of a flood and a heatwave in the same month. In the following paragraph, we rely on Equation 3.2.

Figure 3.3 shows the effects of splitting the regime into three parts. For compound types, the three most significant compound types are shown. In the following paragraphs, many coefficients are shown, yet not all will be discussed in-depth. Instead, patterns will be described. Figure 3.3 shows the effects by hazard type and event type. Two results stand out: first, at the baseline, we find the strongest impact for hail and floods. Figure 3.3 shows that while for hail, single events are indeed the main driver of the results, the negative effects of flood are mainly driven by interactions. This is primarily due to compound flood and temperature events. Second, only for very few hazards are single events themselves significant. Only hail increases the net population change, while storm decreases it. For cumulative events, only temperature anomalies reduce the net population change by -0.8 people per 1000. However, many significant effects are found once accounting for the compounding nature of events. The direction of the coefficients is also sometimes reversed if compound-effects are found. The main takeaway is that spatially compounding events matter and part of the individual effects documented are explained by compound events occurring within a given

year. These compound interactions should be regarded as "new" events, given that, depending on the combination, the direction of the effects on net population change is strongly influenced. A complete list of all interactions can be found in Figure C.6.

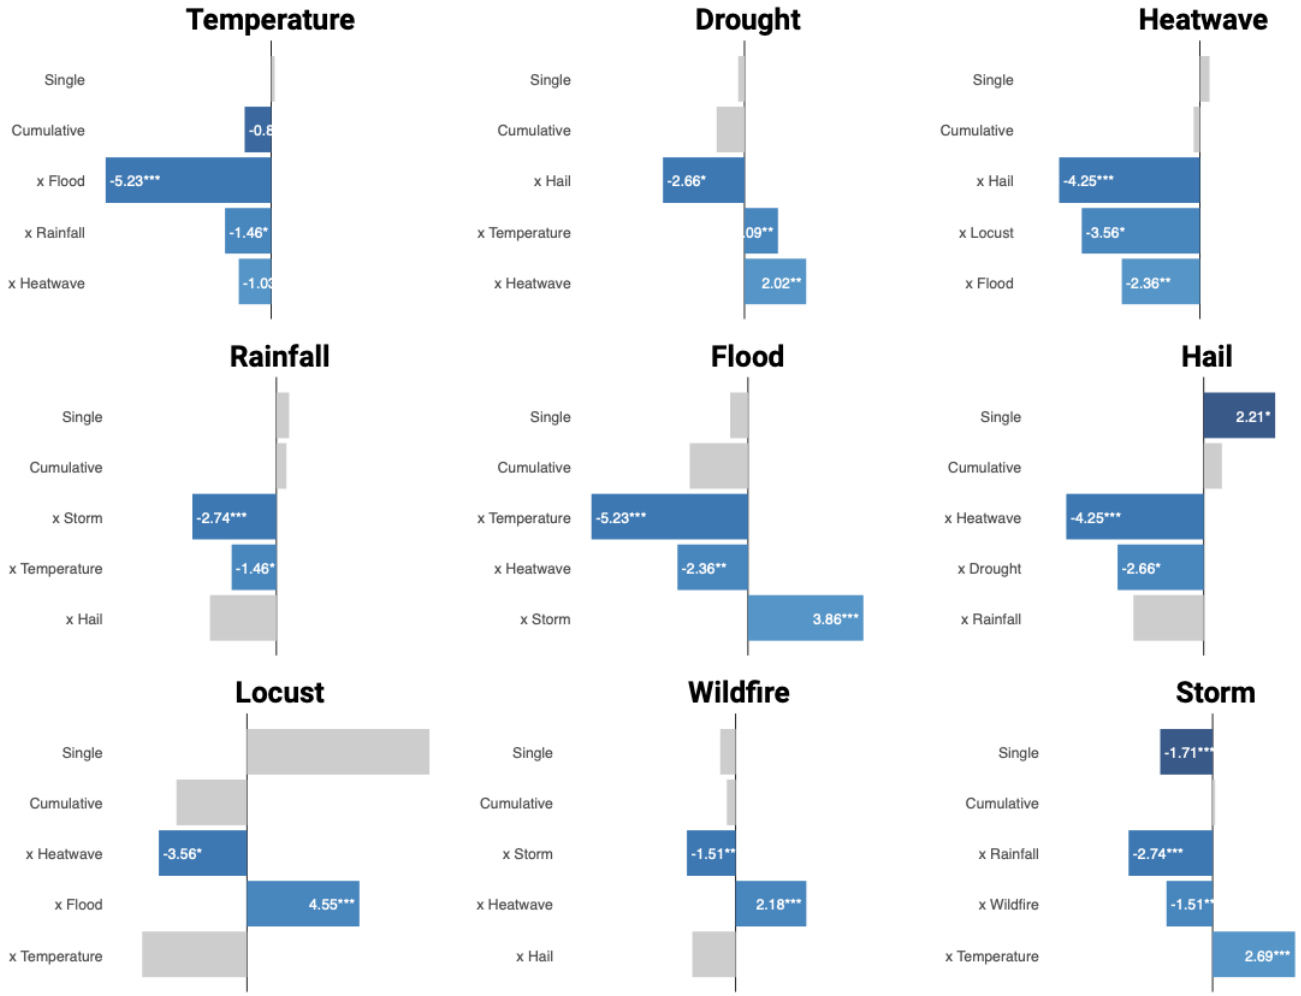

Fig. 3.3. Effects by Hazard Type

Note: *** 0.01, ** 0.05, * 0.1 pval. The figure shows the marginal effects of being exposed to single events, cumulative events (more than one event of the same type within a year) or a compound event (both events occur in the same month). Insignificant results are shown as gray bars. Only the three most significant and largest compound events are shown. All models include district-, year- and continental region-year fixed effects as well as controls for rural population, cropland share and presence of conflict. Standard errors are clustered at the province level.

As in the baseline, the net population change itself is insufficient to explain whether

areas see net losses in population or net gains. Thus, the results of Equation 3.2 are split into net-receiving and net-emitting regions. Figure C.7 splits the results into these two types. At first sight, the number of findings might be overwhelming. In general, more significant effects are found in net negative regions, likely because larger parts of Africa are net emitters rather than net receivers (see Figure C.8 to see which regions are net recipients). Net-receiving areas experience an increased inflow after cumulative floods, cumulative rainfall anomalies, and combinations of storm-temperature and storm-drought events. All other effects either reduce inflow or are not significant at all. This is in contrast to net emitting regions, where cumulative floods lead to increased outflows, but drought- and temperature-storm interactions reduce it. In general, there are more compound types that increase net population losses, likely through migration, in net losing areas. In any case, temperature-flood interactions and drought-storm events show large short-term effects on net population change.

Before proceeding to discuss the robustness and heterogeneity of these findings, it is worthwhile to take a moment to interpret what has been described above. Many climate events, especially compound events, are drivers of negative net population change. In other words, they might drive displacements. The major driver, independent of net-emitting or net-receiving areas, is flood-temperature interactions, as their impact can jointly affect both agricultural income and physical assets. In net-emitting regions, many compound events exhibit an increased loss of population, highlighting underlying displacement pressures, often reflected in patterns of rural-urban migration as households move in search of more secure livelihoods and improved infrastructure. Many of these events involve hydrological events (floods or rainfall). Contrary to the initial analysis for single events only, here, single events rarely matter. In contrast, cumulative sequences of the same type or co-occurrence of different hazards might matter. The presented analysis provides insightful information on the short-term response of the population to different shock regimes and emphasizes the importance of accounting for compounding events or cumulative events that occur within the same year. Future models may also want to study the long-term dynamics of these effects, especially if the recovery from such shocks varies by the type of interaction present.

Robustness of Extended Model

Selection of Variables A key concern with the proposed model is that the inclusion of so many coefficients may lead to imperfect multicollinearity and noisy estimates. To test if the results hold using only the most relevant coefficients, a Post-Double Selection

Lasso estimation is conducted following Belloni, Chernozhukov, and Hansen (2014) and implemented using Ahrens, Hansen, and Schaffer (2024). Hereby, iteratively, only the most relevant hazard types are used, allowing for a "cleaner" comparison. More information is given in subsection C.5. The results are shown in Figure C.16. Overall, the direction and stability of the main coefficients (notably temperature-flood and drought-heatwave) provide evidence that excluding irrelevant coefficients does not substantially change the main findings of the previous section. Changes in magnitude might be attributed to the partial correlation between some hazard classes, which is reduced due to the double selection procedure.

Definition of "Compound" To our knowledge, there does not exist a single clear definition of what "compound" event captures in terms of timing. Papers such as Di Falco et al. (2024) and Petrova et al. (2025) define compound as the cumulative recurrence of the same event over multiple years. In the case of tropical cyclones, wind damage can be compounded by extreme rainfall and flooding within hours or days, as illustrated by Zscheischler, Martius, et al. (2020). We use a rather conservative definition, stating that they must occur in the same month. To compare the effects, we build interaction terms between occurrences within the year and compare the results with the shown model. This shifts the "compound" definition to a "sequencing" definition, where two events occur at any time within the year. The results of the within-year definition are shown in Figure C.18 and discussed in subsection C.5. However, the changed definition to within-year shows a similar direction of coefficients and magnitude in parts smaller than in the baseline. Some events, such as flood-wildfire interactions, become significant once a longer timing interval is allowed between events. In general, the findings affirm that the baseline effects are not strictly driven by the definition of "compound" events. However, changing the definition to an annual level also moves the interpretation away from "compounding" to more sequentially compounding or cumulating events that have different, potentially weaker effects.

3.5 Heterogeneity of the Full Model

So far, the only subsample tested is to study if an area is "net-loosing" or "net-receiving" in terms of population change. However, one might expect different effects also depending on the type of district. In the following, we discuss four different potential sub-samples that might provide insights into where and why the effects described above manifest.

Urban Areas The first heterogeneity tested is by urban versus non-urban areas. Climate change and shocks have been shown to drive urbanization and displacement (e.g., Castells-Quintana, Krause, and McDermott, 2021; Cattaneo and Peri, 2016), while urban areas themselves are increasingly exposed to climate risks (e.g., Rentschler, Avner, et al., 2023). However, due to stronger adaptive infrastructure, more diversified sectoral composition and better access to public services, urban areas may absorb shocks more effectively than rural regions. We expect net population change to be more negative in non-urban areas due to out-migration, whereas urban areas—despite their exposure—may see stable or rising populations,[7] reflecting their role as migration hubs. Columns (1) and (2) of the Table C.3 condition Equation 3.2 into districts with and without major cities to test for heterogeneous effects. The data on cities is derived from the GHS-UCD database (Marí Rivero et al., 2024) and captures dense urban areas. Districts that do not have a city, which are likely more rural areas, show, on average, a stronger magnitude in coefficients (positive or negative), suggesting greater volatility in net population change. For example, the effect of net population change following a single storm is -2.08, almost double the effect of districts that have a city (-1.17). The type of hazards that matter also slightly changes. Although districts with cities experience a strong increase in net population change after temperature-storm (+4.36) and drought-storm (+3.49), those without correspond most strongly to hail-heatwave (-7.24) and heat-wave-locust (-7.23), both of which are potentially damaging to agricultural production. Both types respond negatively to temperature-flood combinations. Similarly to previous findings, single events have small or no effects on net population change, suggesting that some of the effects found in the literature on rural temporary displacement might be largely driven by compound events. The main takeaway is that in both types of regions, compound effects matter; however, rural areas without cities exhibit more volatile, potentially migration-responsive effects. Later, it will be assessed if this might be due to agricultural losses. Districts with cities respond especially to certain compound combinations that often involve temperature-related events and storms, which should be taken into account in adaptation strategies.

Poverty The second heterogeneity tested is poverty. Poorer regions and households may have less absorptive capacity against climate shocks and therefore respond differently. Poor population has been documented to respond differently to climatic shocks (see Cattaneo, Beine, et al., 2019). Following this, the effect on net population change is

[7] Most urban areas show net positive population trends, suggesting they serve as destinations for intra-national migration.

ambiguous: poorer people may have stronger incentives to migrate but fewer financial resources to do so. Given this limited capacity to cope, we expect a reduction in net population—trapping poorer individuals—while wealthier areas may show little to no effect. To compare the results, we use the relative wealth index of Chi et al. (2022) at the district level and split districts in the country of being above or below a country's median in terms of wealth. Columns (3) and (4) of Table C.3 show the results. For poorer regions, the single hazard that matters is hail (+3.99), which increases the net population change. Floods, especially in combination with temperature anomalies and heatwaves, lead to a reduced net population change. Poorer areas show, on average, more negative effects in response to compound events. Richer regions, in contrast, show a higher variability of coefficients. For instance, the range from the smallest to the largest coefficient is seven units (flood temperature to drought storm) in poorer areas, but around 10 units (flood temperature to heat wave-wildfire) in richer areas. Poorer areas, thus, might be more vulnerable to compound events causing effects to be more negative, especially after flood- and heat-related events. Richer areas might have a higher absorptive capacity, thus exhibiting larger positive coefficients, which in turn causes greater variability in coefficients. As discussed above, compound events seem to be more driving results rather than single events.

Climate Vulnerability So far, urban status and relative wealth have been investigated. Next, we assess distinction by climate vulnerability, which is a related, albeit broader concept. Often, climate vulnerability is equated with poverty, which falls short of capturing the full dimension of vulnerability (Doan et al., 2023). The vulnerability of households to climate events may limit their ability to cope with extreme events due to factors such as adaptive capacity, total income, and access to infrastructure (Hallegatte, Vogt-Schilb, Bangalore, et al., 2017; Doan et al., 2023). Thus, climate-vulnerable populations might see disproportionately large effects from extreme events. Hence, we expect the magnitude of the impact of extreme events on the net population to be greater in more climate-vulnerable districts. We test for climate-vulnerable (column (5)) and less climate vulnerable (column (6)) districts using a climate-vulnerability score from 0 to 100. The creation of this score is explained in subsection C.6. In essence, gridded data on wealth, human development, aid dependency, and vulnerable population are standardized and scaled to replicate parts of the INFORM index by the European Joint Research Center (Marin-Ferrer, Poljanšek, and Vernaccini, 2017). It captures, thus, dimensions beyond solely poverty, but also socio-economic vulnerability. Negative net population change is documented in climate-vulnerable areas,

but is not found in low-vulnerability areas. For example, cumulative floods reduce net population change by -6.44 in climate-vulnerable areas but have no effect in more resilient areas. Compound effects continue to exhibit strong effects as before. The strong effects of temperature-flood events are only found in vulnerable areas (-6.61). Only four types of compounds show significant effects and, if so, are mostly positive. The lack of findings of effects in less vulnerable areas might be due to two reasons: either the frequency with which compound events occur is rarer in more resilient areas, or the fact that these areas are more resilient causes effects to be insignificant due to better coping with the shock. In general, vulnerable areas experience a decrease in net population change, potentially leading to increased net population losses, following cumulative and compound events. This emphasizes the limited coping capabilities of these regions. High-resilience areas have mixed effects, but substantially less significant ones, which emphasizes that many of the shocks may be more absorbed by local communities and not necessarily alter short-term population trends.

Remoteness Lastly, assuming that the net population change reflects population mobility, living remotely tends to be more costly for individuals in remote districts due to higher travel costs to urban centers. In such contexts, spatial labor market frictions are likely to intensify with distance, as documented in Brazil following drought events (Albert, Bustos, and Ponticelli, 2021). Therefore, we expect a more negative net population change in remote districts after shocks, reflecting reduced mobility and higher mobility costs. In the final split, we test if areas that are less (column (7)) or more (column (8)) than five hours away from a market town with at least 500 000 people see different effects. The five-hour cut-off point is chosen because it roughly divides our sample in half and is associated with the feasible travel distance that can be covered within a day. The data is based on HarvestChoice; International Food Policy Research Institute (IFPRI) (2016). The main finding is that many of the effects are centered in more remote areas. The effects are especially negative for cumulative and single events. For compound events, again, especially flood-temperature interactions show strong negative effects. In areas closer to major population hubs, single hail events (+3.18) and cumulative floods (+4.08) result in a net population increase, as well as temperature-drought (+1) and drought-storm (+1.86) types. This likely reflects a temporary influx of population or a demographic resilience against compound shocks. Accessibility in general seems to inverse direction of some of the events. In remote areas, we see an increase in net population losses (e.g. through flood-temperature or heatwave-temperature). In contrast, urban-adjacent areas see rather net population gains (especially for cumu-

lative floods). In general, geographic remoteness appears to exacerbate demographic losses following both cumulative and single events. In contrast, spatial proximity to cities seems to correlate with resilience to shocks (fewer significant coefficients), where after some shocks a net population gain is found. It appears that remoteness intensifies vulnerability to shocks.

This section examined patterns in which the effects might vary concerning urban proximity, poverty, climate vulnerability, and remoteness. In general, as previously found, compound events have the strongest and most systematic effects on net population change. For some single-hazard events, effects are observed; however, compound hazards, especially those involving temperature and flood, lead to substantial reductions in net population change in vulnerable, poor, rural, and remote districts. Rural and non-urban districts exhibit significant volatility in effects, highlighting potentially fewer capabilities to absorb incoming shocks, in line with the literature (see for instance Doan et al., 2023; Cattaneo, Beine, et al., 2019). This is further confirmed, as the poorer regions show the most negative effects. The heterogeneity analysis overall reveals that exposure and vulnerability to compound events are spatially unequal and poorer, more rural areas and those with high climate vulnerability tend to be disproportionately affected. The reduction in net population change paired with these areas that exhibit more often net negative population trends is indicative of population losses and displacement, where the local population is likely forced to respond to shocks with relocation. In contrast, the wealthier and more urban parts of Sub-Saharan Africa show greater demographic resilience.

3.6 Mechanisms

Following Cattaneo, Beine, et al. (2019) there are many ways in which climate change could affect net population change and for most valid analysis on true underlying migration household data is required. The channels discussed here will only scratch the surface of the multitude of channels through which extreme events might influence net population change, and notably migratory decisions. Potential channels extend from deteriorated health conditions, loss of agricultural income, to food insecurity, or conflict (Cattaneo, Beine, et al., 2019; Moore and Wesselbaum, 2023; Eberle, Rohner, and Thoenig, 2025). In this section, we focus on the direct impact of extreme events on agricultural pathways and production, as well as food insecurity and income, measured by gross domestic product. Here, it is only tested whether there are effects or

not. In reality, these effects are likely extremely heterogeneous (Burke, Zahid, et al., 2024; Hogan and Schlenker, 2024), as they may vary with sensitivity to the shock, e.g. different crop-combinations respond differently to different shocks, as well as distance to the event. All presented estimations follow Equation 3.2, whereby the dependent variable is replaced with the new variables of interest. For agricultural production and income, as well as addressing food insecurity, instead of using a lagged treatment, the current treatment[8] is used, as its effects may be more immediate than those of population change. The data sources and the revised methodology are described in more detail in subsection C.7.

Agricultural Crop Health Extreme events such as extreme heat could influence the agricultural planting strategy of households (Aragón, Oteiza, and Rud, 2021) and reduce yields. The yield and loss of livestock could lead to an increased migration decision (Blocher, Hoffmann, and Weisz, 2024), affecting the net change of the population. To test this first, we ask whether extreme events affect cropland health positively or negatively in our context. To this end, we compute annual averaged cropland-weighted Normalized Difference Vegetation Index (NDVI) anomalies by district, essentially measuring if over the year cropland in districts becomes "less green". More context for the computation of the measure is provided in subsection C.7. An increase would indicate "improved" plant health, a decrease an "impoverished" plant health[9]. Column (1) of Table C.4 shows nuanced effects with some events increasing and some decreasing vegetation health. Hydrological events, such as rainfall or flood, as well as wildfire seem to increase NDVI anomalies. This is counterintuitive, but could be due to arid regions experiencing prolonged unusual rainfall, where the rain, although potentially destructive in the immediate aftermath, might on an annual level, lead to increased vegetation. In addition, some compound events (hail-rainfall (+0.344) or drought-hail (+0.1557)) could benefit plant health. It is unclear why this is the case, possibly due to similar reasons as for rainfall and floods. More clearly, drought, temperature anomalies and compound events involving drought, locust swarms, and warm and moist environments lead to declines in plant health. Overall, while rainfall-related events can damage crops in the short term on an annual level, they appear to improve NDVI anomalies. The positive effects of wildfires might indicate slash-

[8] In contrast to Equation 3.2 the treatment dummy is not lagged (t instead of t-1).

[9] Note that this analysis is rather descriptive given that NDVI anomalies respond to effects beyond solely hazards. Other background factors, such as soil moisture, soil nutrient density, or farming practices, could have fundamental impacts on NDVI. We believe that the included fixed effects capture part, although by far not all such effects

and-burn practices that might accelerate post-fire recovery. Droughts, especially when compounding over multiple sequences in a year, is a major driver for NDVI losses. When co-occurring with other stressors such as floods or temperature, these combinations might lead to ecological failure and might incentivize net population change. In summary, while certain hydrological and compound events aggregated on an annual level may benefit vegetation health, droughts and their compounds broadly degrade cropland health and may thus influence net population change, along with temperature anomalies.

Agricultural Income Connected to the previous point is the loss of agricultural income. Although changes in NDVI might be indicative of loss of plant health, agricultural production probably accounts for a large portion of rural income[10]. Crop production has been well documented to be affected by extreme events and climate change (Lesk, Rowhani, and Ramankutty, 2016; Lobell, Schlenker, and Costa-Roberts, 2011). Similar to the previous point, losses in crop health and, consequently, agricultural income may translate into significant negative effects on net population change (potentially indicative of out-migration). Thus, we study the effects on staple crops, as they translate either directly into self-subsistence consumption or into local commodities traded on farmers' markets. Following Okou, Spray, and Unsal (2022) the four main staple crops in Sub-Saharan Africa are cassava, wheat, rice and maize, which jointly account for more than half of daily calorie consumption. We obtain sub-national annual crop production statistics from Lee, Anderson, et al. (2025) for these crops and compute logged production by district[11]. Wheat and Rice are omitted from the analysis given that they were frequently missing and had an unstable panel. Usually, in such studies on crop production, the timing of the shock within the crop cycle is also accounted for as shocks at certain stages may result in greater yield losses. Since we use total production instead of yield, the effects of shocks occurring at different points in the crop cycle are assumed to be aggregated into an annual total production estimate. It is essential to note that some of the estimated effects described below may be driven by shocks at specific stages of the crop cycle, a limitation to be considered. Columns (2) and (3) of Table C.4 show the effects on logged production of cassava and maize, respectively. Rainfall, especially if cumulative tends to benefit Cassava, similar to flood-temperature compound events (+0.17 to +0.25% increase in production).

[10] Certainly, production can also be at the commercial, large-scale farming level. However, given that we focus on staple crops, subsistence farming and private commerce of goods is more likely.

[11] If only province-level data is available, they are down-scaled on the district level using information on harvested area by crop from International Food Policy Research Institute (2020)

There appears to be some crop-specific resilience against rainfall or a specific type of compound interaction. However, it is also possible that compound events may not occur during the planting season, thus not significantly influencing the coefficients. Locust swarms, floods and compounding events that involve heatwaves appear destructive, especially for cassava. Cassava, in general, seems to suffer more from compound events, especially with events involving heat and physical shocks (fire or hail). Maize production appears to be more resilient, benefiting even from some compound events that harm cassava. In sum, high-cassava regions may experience more substantial losses in agricultural income, especially through locust swarms (in the Sahel region), floods, and heat-related shocks that may influence net population change. In contrast, regions with dominant maize production may, in parts, even benefit from certain types of compound events. However, note that these are trends based on aggregated production numbers. We did not compute the crop-specific response for each level of hazard intensity (i.e. dose-response function), which would have provided insights into the actual production patterns using models such as those mentioned in Blanc and Schlenker (2017). Future research will need to extend these studies by using nonlinear impact models and accounting for irrigation practices, crop cycle shifts, and local farming practices to fully account for crop responses to compound shocks.

Food Insecurity Next, we discuss two measures of economic fragility, notably food insecurity and loss of economic growth. Food insecurity can trap households in a lack of options to migrate (see, e.g., Sadiddin et al., 2019) or leave them with no financial assets to migrate, resulting in stagnation of population change. In this section, we examine the hazards that contribute to food insecurity, as measured by the Integrated Food Security Phase Classification (IPC) of the Famine Early Warning Systems Network (Andrée et al., 2020). The data covers 18 African countries from 2009 until the end of our study period. Positive effects indicate increases in food insecurity, so worse conditions. Column (4) of Table C.4 shows the results. Single hazards that exhibit significant effects are temperature (+0.0548), drought (+0.0751), and hail (+0.0757), all of which increase food insecurity, likely due to their impact on agricultural production. Cumulative rainfall reduces food insecurity (-0.11) likely because persistent rainfall anomalies in arid regions might create agricultural stability. Cumulative droughts and temperature, in reverse, show a strong increase in food insecurity. Some compound events show a reduction in food insecurity, notably hail-temperature (-0.24), drought-locust (-0.45) and hail-wildfire (-0.5). This is surprising and may be explained by either a local adaptive response to these shocks or national or international intervention that supported re-

gions affected by such shocks. Flood-locust (0.35) and drought-rainfall (0.18) exhibit the strongest compound response, contributing to increased food insecurity. In terms of net population change, these results suggest that cumulative events, particularly drought and temperature, may impact net population change through food insecurity. Some compound interactions, such as flood-locust or drought-rainfall, may create immediate and acute effects, potentially also influencing net population change if they overwhelm local coping capacity.

Economic Growth Finally, economic crisis and deprivation could provide incentives to migrate if the disparities between regions grow. Thus, as a final mechanism, the effect on GDP growth is assessed. The effects of disasters and extreme events are well-researched in the literature (see, e.g., Felbermayr, Gröschl, et al., 2022; Hsiang and Jina, 2014; Cevik and Jalles, 2024). We expect little to no effect on GDP growth given that the estimates on GDP are likely highly-elastic and overweight urban areas due to the usage of nightlight for calibration (Gibson, Kim, and Li, 2024). Thus, we expect little to no effects on GDP growth. Here, we rely on province-level estimates from Zhang et al. (2024), which distribute GDP using a deep-learning model based on nightlights. Estimates are downscaled to the district level using population shares. More context is given in subsection C.7. Column (5) of Table C.4 shows broadly weak or nonsignificant effects. Single and cumulative events have a minimal impact on GDP growth. Flood and cumulative temperature reduce it marginally. Certain compound types like temperature-storm (-0.0061) and wildfire-storm (-0.0076) show stronger effects. The strongest effects are found for locust-rainfall events (-0.02). Surprisingly, cumulative hail and drought-flood events accelerate economic growth, potentially due to quick recovery or sectoral shifts. In summary, the effects of extreme events on district GDP growth are limited, as tested here. Compound events, particularly those involving climate-biological interactions, appear to have the most negative impacts. The lack of impact dimension is likely due to GDP estimates being mainly based on downscaled nightlight, where large parts of rural Africa show little to no nightlight. This is a common shortcoming for applications involving nightlight in rural regions of under-electrified regions. As such, the findings might obfuscate some of the actual impacts of extreme events on GDP. Nonetheless, we believe that these effects are indicative of underlying behavior on a fairly local scale, as large parts of the sample are rural and are likely to be minimally impacted by events in terms of nightlight.

This section outlines four mechanisms that might explain why extreme events impact net population change in Sub-Saharan Africa. Extreme events have heterogeneous im-

pacts on agricultural health, with both single and compound events - especially temperature anomalies and droughts - influencing cropland health. Drought-temperature combinations might influence net population change. This is in line with the reduced net outflow in emitting areas found in Figure C.7 that might be caused by loss of income. Crop types also influence vulnerability, as cassava appears to be more vulnerable compared to maize. Heat-related and physical events (such as wildfires) are especially damaging, suggesting the importance of crop composition. Drought and heat are also associated with higher food insecurity. This is also true for certain compound types, such as flood-locust and drought-rainfall. Food insecurity itself might force displacement or trap the population. Cumulative droughts significantly increase net population loss in districts already in decline (see Figure C.7). The impacts of GDP are minimal, probably due to the data limitations of nightlight estimates in rural Africa. Overall agricultural disruption, income shocks, and food insecurity might explain part of the baseline effects on net population change, particularly in poorer, more rural, and remote districts as discussed before. Interestingly, the flood-temperature compound event appears to increase crop production but has no other notable effects. The absence of impact, combined with significant effects in the baseline, suggests that other channels, such as physical destruction or agricultural losses beyond cassava and maize, may cause the baseline effects.

3.7 Conclusion

This paper examines the short-term effects of extreme events and their temporal overlaps, known as "compound" events, on the net change in population in sub-Saharan Africa between 2000 and 2019. Relying on a novel dataset that aggregates nine common hazards and combines them with data on the net population trend at the district level, we reveal a complex pattern of impact on the local population. Not accounting for the complexity of events and adopting a more "standard" binary definition of having an event or not, we find that floods lead to a reduction, notably outflow, of local population, while hail tends to increase population in net-receiving areas. Splitting the treatment, however, into three treatment regimes within each year: single events, repeated cumulative events, and "compound events", we document a wide range of effects. Many of the previous insignificant effects are, in fact, significant, once accounting for overlapping occurrences of events. Single events in place of storms rarely matter. Instead, different combinations of hazard events might lead to fundamentally different

impacts. Most prominent are temperature-flood interactions causing displacement[12], contrary to heatwave-wildfire and flood-storm that see reduced outflows of district population. The first major finding is that rather than considering individual hazard classes, accounting for intraannual overlap of events adds important insights. Hereby, each specific hazard-combination should be regarded separately due to its potentially different impact channels.

The effects found are centered on population changes in vulnerable, poor, rural and remote districts of Sub-Saharan Africa. However, there is high volatility in the effects, indicating a low capacity to cope with the shock. Given that these regions, on average, see net population losses, it highlights strong displacement effects of particularly compound events. Wealthier, more urban parts show, in contrast, more resilience against shocks. The mechanism part provided insight that compound events have a significant impact on agricultural health and production, as well as food insecurity, primarily through heat and drought-related events. However, we also discuss the possibility that other mechanisms might be at play, such as the loss of physical assets. Thus, the second main finding is that not only do compound and cumulative events in a year matter, but also they disproportionately affect agricultural-dependent rural areas that tend to be poorer and more remote.

The final takeaway from this article is methodological in nature. This article does not claim causality because net population data is strongly dependent on interpolated census data. Additionally, we document that our panel model, with its three sets of fixed effects, fails to capture spatial spillovers, temporal dynamics, and within-month nonlinearities. While accounting for fixed effects and including other climatic background stressors to mitigate a multitude of potential biases, we document, at least for some hazards, that spillovers and nonlinearities exist. While seemingly a limitation, we argue that these observed findings also contribute to the literature by documenting that, in contrast to most currently tested models, empirical models must account for all three conditions: first, the models need to incorporate a spatial treatment dimension that allows spillovers between spatial units. Second, it needs to account for temporal dynamics. To our knowledge, the current difference-in-difference literature (e.g. Baker et al., 2025) does not allow for treatment reversal under the assumption that each treatment-combination in itself carries over a different number of years (depending on the socio-economic context of the place). For disaster event studies, excluding this makes the

[12] Net population loss in net-emitting and reduced net population gain in net-receiving districts

model rely on strong econometric assumptions. Third, the estimator needs to account for nonlinearities, particularly for interactions of different hazard-pairs. Binary joint occurrence, as done in this article, fails to capture the two respective intensity levels that themselves might influence economic outcomes. To our knowledge, most studies on impacts on socioeconomic outcomes focus on one or maximal two of these concerns, but not all three.

The findings have three main policy implications. First, findings that compound events might elicit stronger responses should incentivize policymakers to shift from the notion that hazards are individual events to recognizing their interconnectedness within a complex system. The tested compound event has distinct and stronger impacts and should be regarded as a "new" type of event. The integration of compound event forecasting should be increasingly added to early warning systems and disaster risk reduction protocols. Second, the most significant effects are found in poor, remote and rural areas and in parts through the agricultural channel. The findings support the idea of agricultural support programs and also shift attention to social protection in these areas. Examples could be compound-resistant crop systems or improved community safety nets. Third and finally, there should be an increased effort and push in the development of an advanced impact model that accounts for spatial dependency, temporal carry-over of effects, as well as nonlinearities simultaneously, in combination with high-granular and high-frequency data, to fully uncover the socioeconomic response to this form of extreme events.

This study has a limited scope and extent. As mentioned above, it does not fully address all methodological concerns present in the literature. At the same time, it does not capture the full complexity of the hazard impacts, although more than is usually shown. First, given data limitations, overlapping events are defined here as those that occur in the same month. However, within the month, it may also matter, especially for rapid-onset events such as floods or hail, how many days separate the two events. With the data used here, this is not possible. Second, in contrast to Di Falco et al. (2024) and Petrova et al. (2025) "compound events" are defined within-year not across years. Both are equally important and deserve attention. Third, in contrast to Thalheimer, Choquette-Levy, and Garip (2022), we do not include the structural and social vulnerability as a driver for compound events, which is equally important. Moreover, "compound" events are defined by bivariate interactions of events; however, interactions of more than two event types are also possible. Lastly, we do not account for the order of events. However, the order in which events occur may change the direction

and magnitude of the coefficient. For example, floods following drought may be more severe if the soil is dry. In future research, a more holistic definition of "compound" should be explored, encompassing both temporal and social dimensions, as well as the sequence of events. Our study focuses on net population change rather than migration. Thus, similar studies should be repeated that combine high-resolution hazard intensities with individual-level household information on migratory decisions. Finally, future work should carefully disentangle the underlying personal motivation for people to move, which is impossible with district-aggregate-level data, yet this concern has already been raised many times in the literature (e.g. Beine and Jeusette, 2021; Hoffmann, Šedová, and Vinke, 2021).

Studying the interdependence of climate and population has a rich history in the fields of development economics and demography. Advances in statistical models and big data enable researchers to provide a more in-depth analysis of population and migration dynamics, particularly in rural regions, where household surveys may be less frequent. With global warming that influences not only the variability and intensity of individual events but also the co-occurrence of different events, it becomes more important than ever to study how these "new" events affect local communities. Failing to grasp these dynamics can lead policymakers to misdirect adaptation efforts, focusing on areas experiencing a rise in a single-hazard class rather than recognizing the compounding effects of multiple hazards in other regions, exacerbating the impact of climate change on rural populations by misguided policies.

C Appendix to Chapter 3

C.1 Additional Tables

Table C.1
Thresholds for Cut-off Intensity (Reinhardt, 2024a)

Hazard	Measure	Onset Type	Dimension	Share of sample above threshold
Temperature (°C)	z-score	Slow	Regional	3.8%
Heatwave (WBGT)	5-day MA	Slow	Regional	2.2%
Precipitation (mm)	z-score	Rapid	Regional	1.1%
Drought	SPEI-3 & SPEI-6	Slow	Regional	2.8%
Tropical Storm	m/s	Rapid	Local	<1%
Wind gusts	m/s	Rapid	Local	<1%
Flood	DFO Index	Rapid	Local	1.3%
Hail	Days	Rapid	Local	< 1%
Forest Fire	Share Burned	Slow	Local	2.7%
Locust Swarms[x]	Hopper Count	Slow	Regional	<1%

Note:[x] is extended to Reinhardt (2024a) from Food and Agriculture Organization (FAO) of the United Nations (2023). WBGT: Wet Bulb Globe Temperature, MA: Moving Average, SPEI: Standardized Precipitation Evapotranspiration Index, DFO: Dartmouth Flood Observatory. Onset type explains the time it needs for these events to show substantial effects relative to the other included hazards. Original input sources based on: Muñoz Sabater, 2019; Mistry, 2020; Gebrechorkos et al., 2023; Done et al., 2020; Knapp et al., 2010; Copernicus Climate Change Service, 2024a; Brakenbridge and Dartmouth Flood Observatory, 2023; Prein and Holland, 2018b; Copernicus Climate Change Service, 2019; Vermote and Wolfe, 2015.

Table C.2
Overview on Three Treatment Types

Event	Observations	Districts	Years	Event	Observations	Districts	Years
Single Events				**Compound Events**			
Temperature	15,917	3,821	20	Drought x Flood	244	227	14
Drought	6,586	3,192	20	Drought x Hail	93	87	10
Heatwave	6,906	2,066	19	Drought x Heatwave	1,189	941	16
Rainfall	9,667	3,577	20	Drought x Locust	70	54	16
Flood	6,926	2,887	20	Drought x Rainfall	152	133	17
Hail	1,315	410	16	Drought x Temperature	5,335	2,817	20
Locust	37	26	16†	Drought x Wildfire	1,888	1,138	19
Wildfire	8,612	1,673	19	Drought x Storm	131	74	14
Storm	439	157	20	Flood x Hail	98	75	6
Cumulative Events				Flood x Heatwave	675	506	11
Temperature	8,599	2,990	20	Flood x Locust	45	37	9
Drought	2,460	1,655	20	Flood x Rainfall	815	724	18
Heatwave	4,885	1,917	20	Flood x Temperature	216	210	15
Rainfall	2,105	1,420	20	Flood x Wildfire	419	309	17
Flood	109	90	20	Flood x Storm	78	62	10
Hail	345	98	20	Hail x Heatwave	134	113	4
Locust	127	51	20	Hail x Locust	20	13	8†
Wildfire	8,325	1,229	20	Hail x Rainfall	66	57	12
Storm	365	51	20	Hail x Temperature	215	149	13
				Hail x Wildfire	115	72	13
				Hail x Storm	12	4	7†
				Heatwave x Locust	71	52	10
				Heatwave x Rainfall	983	744	14
				Heatwave x Temperature	1,594	1,064	16
				Heatwave x Wildfire	1,835	715	12
				Heatwave x Storm	32	28	4†
				Locust x Rainfall	105	61	18
				Locust x Temperature	109	63	15
				Locust x Wildfire	6	6	3†
				Locust x Storm	1	1	1†
				Rainfall x Temperature	1,047	833	20
				Rainfall x Wildfire	2,273	1,072	19
				Rainfall x Storm	105	69	19
				Temperature x Wildfire	4,694	1,698	19
				Temperature x Storm	80	46	16
				Wildfire x Storm	173	37	19

Note: The table plots for each event the total number of districts that received a treatment of that kind over the 20-year period (observations). It also shows the total number of districts that receive ever this kind of treatment (Districts) and in how many years this treatment occurs (Years). The † indicates samples that in Equation 3.2 are omitted, given that either fewer than 30 districts or fewer than 30 observations are recorded.

Table C.3
Heterogeneity: By Stratifiers

	Has City?		Intra-national wealth is...		Climate Vulnerability		Travel Time is..	
	No	Yes	Below p50	Above p50	More	Less	<5h	>5h
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
Single Events								
Temperature (TP)	0.42	-0.29	-0.16	0.43	-0.18	0.36	-0.04	0.38
	(0.75)	(-0.92)	(-0.46)	(0.69)	(-0.51)	(0.61)	(-0.13)	(0.61)
Drought (DR)	-0.36	-0.03	-0.63*	0.17	-0.43	0.24	-0.26	0.04
	(-0.73)	(-0.08)	(-1.66)	(0.37)	(-1.13)	(0.46)	(-0.63)	(0.09)
Heatwave (HW)	0.59	-0.05	0.15	0.08	0.65	-0.08	-0.17	1.09
	(0.92)	(-0.10)	(0.27)	(0.10)	(0.97)	(-0.14)	(-0.51)	(1.21)
Rainfall (RF)	0.22	0.55	0.00	0.69*	-0.15	0.58	0.22	0.51
	(0.48)	(1.53)	(0.01)	(1.86)	(-0.48)	(1.35)	(0.53)	(1.24)
Flood (FL)	-0.90*	-0.44	-0.55	-0.60	-0.57	0.04	0.08	-1.28**
	(-1.68)	(-0.96)	(-1.17)	(-1.03)	(-1.01)	(0.08)	(0.17)	(-2.12)
Hail (HL)	4.02**	0.22	3.99**	-0.53	2.57*	0.69	3.18*	0.75
	(2.33)	(0.27)	(2.30)	(-0.53)	(1.77)	(0.42)	(1.83)	(0.70)
Locust (LO)	10.55	2.13	9.27	4.47	8.35	6.20	17.90	2.33
	(1.59)	(0.34)	(1.15)	(0.69)	(1.23)	(0.84)	(1.36)	(0.52)
Wildfire (WF)	-0.65	-0.21	0.12	-1.39**	-0.46	-0.18	-0.03	-0.66
	(-1.39)	(-0.40)	(0.26)	(-2.25)	(-0.84)	(-0.37)	(-0.06)	(-1.21)
Storm (WI)	-2.08***	-1.17*	-1.59**	-1.37*	-2.95***	-0.60	-1.13	-1.98***
	(-2.85)	(-1.83)	(-1.97)	(-1.91)	(-2.96)	(-0.95)	(-1.10)	(-2.82)
Cumulative Events								
Temperature (TP)	-0.88	-0.81	-1.84***	0.17	-1.60***	0.60	0.14	-1.15*
	(-1.33)	(-1.62)	(-2.82)	(0.27)	(-3.00)	(1.03)	(0.25)	(-1.78)
Drought (DR)	-1.09	-0.77	-2.01***	0.26	-1.22*	0.54	-2.40**	0.24
	(-0.99)	(-1.38)	(-2.70)	(0.23)	(-1.85)	(0.43)	(-2.53)	(0.27)
Heatwave (HW)	-0.75	0.12	-0.30	-1.03	1.69**	-1.58	-0.86	0.83
	(-0.94)	(0.16)	(-0.55)	(-0.95)	(2.25)	(-1.43)	(-1.37)	(0.94)
Rainfall (RF)	0.10	0.48	-0.43	0.99	0.13	1.17*	-1.27	0.74
	(0.17)	(0.78)	(-0.60)	(1.55)	(0.18)	(1.81)	(-1.21)	(1.26)
Flood (FL)	-1.15	-3.01	0.24	-4.32	-6.88**	1.88	4.08**	-5.02**
	(-0.63)	(-0.81)	(0.12)	(-1.37)	(-2.29)	(1.41)	(2.16)	(-2.14)
Hail (HL)	0.16	0.80	2.27	-3.17	0.48	0.70	2.00	-5.78***
	(0.03)	(0.56)	(0.74)	(-1.41)	(0.18)	(0.32)	(0.76)	(-4.05)
Locust (LO)	-3.44	-2.19	-0.34	-4.28	0.63	-5.26	-15.18	-0.42
	(-0.68)	(-1.38)	(-0.18)	(-0.93)	(0.66)	(-0.66)	(-0.91)	(-0.34)
Wildfire (WF)	-0.41	-0.29	0.24	-1.00	-0.28	-0.25	-0.11	-0.54
	(-0.75)	(-0.53)	(0.53)	(-1.15)	(-0.52)	(-0.41)	(-0.17)	(-0.95)
Storm (WI)	-0.74	0.80	-0.20	0.44	-3.98*	0.03	0.82	-1.35
	(-0.58)	(0.51)	(-0.16)	(0.35)	(-1.91)	(0.04)	(1.21)	(-1.03)
Compound Events								
DR x FL	-1.61	-0.05	-0.70	-0.14	-0.53	-0.72	0.37	-0.39
	(-0.64)	(-0.05)	(-0.67)	(-0.04)	(-0.37)	(-0.52)	(0.33)	(-0.19)
DR x HL	-0.82	-3.48**	-3.49*	-2.05			-3.74*	-0.75
	(-0.24)	(-2.53)	(-1.66)	(-0.89)			(-1.94)	(-0.33)
DR x HW	3.41**	1.32	0.61	2.94**	2.10*	1.60	0.52	3.48**
	(2.25)	(1.39)	(0.64)	(2.04)	(1.91)	(1.11)	(0.61)	(2.38)
DR x LO	0.03	1.28						

Continued on next page

	Has City?		Intra-national wealth is...		Climate Vulnerability		Travel Time is..	
	No	Yes	Below p50	Above p50	More	Less	<5h	>5h
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
	(0.01)	(0.60)						
DR x RF	-3.07	1.02	0.56	-0.76	-0.73	0.01		
	(-1.59)	(0.58)	(0.40)	(-0.43)	(-0.46)	(0.00)		
DR x TP	1.32*	0.65	1.55**	0.71	1.21**	0.58	1.00*	0.81
	(1.96)	(1.59)	(2.41)	(1.39)	(2.14)	(1.19)	(1.74)	(1.36)
DR x WF	-1.09	-0.14	-1.50*	0.41	-1.06	1.49	-1.35	-0.64
	(-0.92)	(-0.19)	(-1.88)	(0.36)	(-1.30)	(1.36)	(-1.23)	(-0.69)
DR x WI	3.80***	3.49***	2.99***	3.37***	2.33*	0.73	1.86*	4.47***
	(2.87)	(3.03)	(2.74)	(2.62)	(1.93)	(0.97)	(1.66)	(3.50)
FL x HL	2.84**	1.27						
	(2.09)	(1.63)						
FL x HW	-2.32*	-2.19**	-1.92**	-4.55*	-0.88	-3.96	-2.13**	-1.95
	(-1.79)	(-2.32)	(-2.30)	(-1.66)	(-0.78)	(-1.60)	(-2.13)	(-1.23)
FL x RF	-0.07	-1.99	0.30	-2.24	0.65	-2.83	-1.11	-0.51
	(-0.08)	(-1.13)	(0.42)	(-1.15)	(0.86)	(-1.33)	(-0.58)	(-0.54)
FL x TP	-5.85*	-4.28***	-4.64*	-5.51**	-6.61**	-2.49	-1.50	-8.30***
	(-1.92)	(-2.70)	(-1.72)	(-1.98)	(-2.57)	(-1.08)	(-0.62)	(-3.43)
FL x WF	-0.21	3.69	1.58	1.15	1.67	-0.89	-0.21	1.58
	(-0.09)	(1.48)	(1.50)	(0.38)	(0.68)	(-0.84)	(-0.12)	(0.70)
FL x WI	4.95***	2.04	1.58	6.07	8.56***	1.48		
	(2.91)	(0.80)	(1.59)	(1.64)	(3.45)	(1.27)		
HL x HW	-7.24**	-2.78**						
	(-2.56)	(-2.41)						
HL x RF	9.15*	-8.22						
	(1.77)	(-0.87)						
HL x TP	0.95	2.51**	0.59	2.94**	-0.65	3.45*	0.80	1.54
	(0.18)	(2.30)	(0.14)	(2.55)	(-0.18)	(1.94)	(0.21)	(1.18)
HL x WF	-1.99	-0.79	-1.57	-1.73				
	(-1.01)	(-0.79)	(-1.14)	(-0.89)				
HW x LO	-7.23**	2.62						
	(-2.23)	(1.49)						
HW x RF	-0.66	-0.79	-1.06*	0.05	-1.08*	-0.51	-1.37	-0.15
	(-0.70)	(-1.16)	(-1.75)	(0.04)	(-1.80)	(-0.44)	(-1.40)	(-0.20)
HW x TP	-1.92*	-0.16	-0.60	-1.90*	-0.69	-2.08*	-0.71	-1.91**
	(-1.82)	(-0.30)	(-0.83)	(-1.93)	(-1.21)	(-1.69)	(-0.68)	(-2.36)
HW x WF	2.86***	1.44**	0.51	5.41***	2.79***	1.20	0.60	3.02***
	(2.61)	(2.11)	(1.23)	(3.14)	(2.91)	(1.21)	(1.00)	(2.66)
LO x RF	-2.43*	1.41			0.08	-1.71		
	(-1.68)	(0.71)			(0.07)	(-0.59)		
LO x TP	-6.54	-0.17			-0.85	-8.93		
	(-0.53)	(-0.11)			(-0.48)	(-0.45)		
RF x TP	-2.51*	0.11	-0.27	-2.67*	-1.12	-1.20	-2.46	-1.07
	(-1.81)	(0.15)	(-0.31)	(-1.83)	(-1.20)	(-0.71)	(-1.60)	(-1.05)
RF x WF	0.12	0.41	0.83*	-0.40	0.57	0.36	0.77	-0.07
	(0.13)	(0.76)	(1.90)	(-0.36)	(0.73)	(0.48)	(0.89)	(-0.10)
RF x WI	-1.79	-4.22***	-2.56***	-2.86**	-0.01	-1.47***		
	(-1.55)	(-2.62)	(-2.83)	(-2.21)	(-0.00)	(-2.84)		
TP x WF	-0.76	-0.43	-0.03	-1.18	-0.52	-0.51	0.64	-0.98
	(-1.12)	(-0.78)	(-0.06)	(-1.42)	(-0.88)	(-0.75)	(1.13)	(-1.55)
TP x WI	2.25**	4.36***	2.41	2.88***	0.11	1.35*		

Continued on next page

	Has City?		Intra-national wealth is...		Climate Vulnerability		Travel Time is..	
	No	Yes	Below p50	Above p50	More	Less	<5h	>5h
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
	(2.44)	(3.08)	(1.54)	(3.13)	(0.05)	(1.70)		
WF x WI	-1.89**	-1.76***	-2.55**	-1.05			-0.46	-1.86
	(-2.00)	(-2.80)	(-2.10)	(-1.28)			(-0.92)	(-1.28)
R^2	0.428	0.572	0.514	0.468	0.659	0.322	0.432	0.517
Obs.	38,448	31,068	34,758	34,758	34,758	34,758	34,380	35,082
Net Population Change	-4.22	1.23	-3.00	-0.60	-3.46	-0.11	-1.82	-1.87
Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Region× Year FE	✓	✓	✓	✓	✓	✓	✓	✓

Note: p-values: *<0.1,**<0.05,***<0.01. The table shows the results of Equation 3.2. Only events are included where, in the subsample, at least 30 events have been recorded to avoid outlier events. All models are clustered at the province-year level.

Table C.4
Mechanisms: Plant Health, Production, and Economic Fragility

	Plant Health	Log(Production)		Economic Fragility	
	NDVI Anomaly	Cassava	Maize	IPC Score	GDP-Growth
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Single Events					
Temperature (TP)	-0.06***	0.01	-0.04**	0.05***	-0.00
	(-6.85)	(0.44)	(-2.27)	(2.83)	(-0.41)
Drought (DR)	-0.09***	-0.07**	0.01	0.08***	-0.00
	(-6.82)	(-2.58)	(0.27)	(3.33)	(-0.20)
Heatwave (HW)	0.01	0.03	0.02	-0.03	0.00
	(0.98)	(0.81)	(0.45)	(-0.66)	(0.20)
Rainfall (RF)	0.06***	0.06***	0.02	-0.03	-0.00
	(6.96)	(2.70)	(0.90)	(-1.50)	(-1.37)
Flood (FL)	0.10***	0.03	-0.05*	-0.02	-0.00**
	(6.33)	(0.61)	(-1.78)	(-0.78)	(-2.35)
Hail (HL)	-0.07***	-0.20***	0.03	0.08*	0.00
	(-2.87)	(-2.95)	(0.38)	(1.73)	(0.76)
Locust (LO)	-0.01	-0.51***	-0.10	-0.04	0.01
	(-0.13)	(-10.95)	(-0.17)	(-0.48)	(0.78)
Wildfire (WF)	0.06***	0.06**	0.03	0.04	-0.00
	(4.39)	(2.27)	(1.41)	(1.13)	(-0.15)
Storm (WI)	0.08	-0.14**	0.10*	0.13	0.00
	(1.56)	(-2.09)	(1.68)	(0.99)	(0.56)
Cumulative Events					
Temperature (TP)	-0.08***	-0.01	-0.01	0.07***	-0.00***
	(-6.31)	(-0.42)	(-0.21)	(2.65)	(-2.97)
Drought (DR)	-0.21***	-0.01	0.00	0.17***	-0.00
	(-8.69)	(-0.27)	(0.05)	(4.05)	(-0.42)
Heatwave (HW)	-0.02	-0.18***	-0.00	-0.09	-0.00
	(-0.94)	(-2.88)	(-0.08)	(-1.63)	(-0.29)

Continued on next page

	Plant Health	Log(Production)		Economic Fragility	
	NDVI Anomaly	Cassava	Maize	IPC Score	GDP-Growth
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Rainfall (RF)	0.10***	0.08**	-0.02	-0.12***	0.00
	(5.88)	(2.01)	(-0.64)	(-3.37)	(0.43)
Flood (FL)	0.26***	-0.51***	-0.09	0.07	0.00
	(3.13)	(-2.88)	(-1.30)	(0.58)	(0.45)
Hail (HL)	-0.07	-0.08	-0.01	0.04	0.01*
	(-1.51)	(-1.31)	(-0.04)	(0.29)	(1.85)
Locust (LO)	0.16**	-0.78	-0.37	-0.20	0.02
	(2.07)	(-1.51)	(-1.60)	(-1.28)	(0.79)
Wildfire (WF)	0.09***	0.11***	0.04	0.08	0.00
	(4.79)	(2.91)	(1.35)	(1.29)	(0.30)
Storm (WI)	0.12	-0.05	0.02	0.20	-0.00
	(1.12)	(-0.45)	(0.51)	(0.66)	(-1.17)
Compound Events					
DR x FL	-0.31***	0.09	-0.01	-0.03	0.00**
	(-5.14)	(0.67)	(-0.13)	(-0.38)	(2.45)
DR x HL	0.16**	-0.20	-0.05	0.18	-0.00
	(2.22)	(-0.81)	(-0.37)	(1.35)	(-0.13)
DR x HW	-0.03	-0.25***	-0.13	-0.23***	-0.00
	(-0.92)	(-3.76)	(-1.57)	(-3.63)	(-0.75)
DR x LO	0.11	0.00	-0.03	-0.46***	-0.01
	(1.50)	(.)	(-0.15)	(-3.50)	(-0.79)
DR x RF	-0.00	-0.05	-0.17	0.18*	-0.01
	(-0.04)	(-0.36)	(-0.82)	(1.72)	(-0.66)
DR x TP	-0.15***	-0.03	-0.08**	0.13***	-0.00
	(-8.43)	(-0.92)	(-2.18)	(4.55)	(-0.53)
DR x WF	-0.03	-0.06	-0.02	-0.10**	0.00
	(-1.28)	(-1.27)	(-0.35)	(-2.23)	(1.02)
DR x WI	-0.08	-0.04	-0.17***	0.04	-0.00
	(-0.59)	(-0.34)	(-3.77)	(0.28)	(-0.28)
FL x HL	0.02	-0.17	0.11	0.04	-0.01
	(0.34)	(-1.21)	(0.22)	(0.47)	(-1.33)
FL x HW	-0.08**	-0.06	-0.02	0.14*	0.00
	(-2.12)	(-0.82)	(-0.36)	(1.77)	(0.24)
FL x LO	0.05	0.00	0.02	0.36***	-0.01
	(0.50)	(.)	(0.07)	(3.74)	(-1.29)
FL x RF	0.04	-0.14	-0.00	-0.01	0.00
	(0.95)	(-1.46)	(-0.01)	(-0.25)	(0.18)
FL x TP	0.03	0.17*	0.46***	-0.08	0.00
	(0.62)	(1.73)	(4.73)	(-0.77)	(0.22)
FL x WF	-0.13***	0.11	-0.10	0.01	-0.00
	(-2.96)	(1.31)	(-1.03)	(0.25)	(-1.56)
FL x WI	-0.02	0.19	-0.20**	0.12	0.01
	(-0.15)	(1.49)	(-2.38)	(0.88)	(1.59)
HL x HW	-0.07	-0.30**	-0.50**	-0.29*	0.00
	(-1.20)	(-2.57)	(-2.03)	(-1.75)	(0.56)
HL x RF	0.34***	0.09	-0.09	-0.11	-0.00
	(6.60)	(0.61)	(-0.98)	(-0.87)	(-0.46)
HL x TP	0.01	0.03	-0.05	-0.25***	0.00
	(0.17)	(0.52)	(-0.97)	(-2.83)	(0.71)
HL x WF	0.09*	0.17	-0.14	-0.50***	0.00
	(1.94)	(0.81)	(-1.25)	(-4.35)	(0.42)
HW x LO	0.06	0.00	-0.07	0.05	0.00

Continued on next page

	Plant Health	Log(Production)		Economic Fragility	
	NDVI Anomaly	Cassava	Maize	IPC Score	GDP-Growth
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
	(0.79)	(.)	(-0.15)	(0.33)	(0.44)
HW x RF	-0.01	-0.03	-0.00	0.00	-0.00
	(-0.30)	(-0.46)	(-0.02)	(0.03)	(-1.49)
HW x TP	0.07***	-0.02	-0.12*	-0.03	0.00
	(3.97)	(-0.56)	(-1.72)	(-0.51)	(1.27)
HW x WF	-0.01	0.07	-0.01	-0.25***	-0.00
	(-0.67)	(1.61)	(-0.28)	(-4.14)	(-0.72)
LO x RF	-0.11**	0.00	0.51**	0.17	-0.02**
	(-2.58)	(.)	(2.47)	(1.34)	(-2.02)
LO x TP	-0.00	0.05	0.35**	0.11	0.00
	(-0.02)	(0.83)	(2.24)	(1.62)	(0.32)
RF x TP	-0.06***	-0.10**	0.05	0.10**	0.00
	(-3.33)	(-1.97)	(1.22)	(2.15)	(0.89)
RF x WF	0.03**	0.04*	-0.04*	0.06	-0.00
	(1.99)	(1.65)	(-1.77)	(1.35)	(-0.71)
RF x WI	0.10	-0.14	-0.03	0.27	-0.00
	(1.09)	(-1.06)	(-0.28)	(1.43)	(-0.17)
TP x WF	-0.01	0.07**	0.04	0.06*	-0.00*
	(-1.00)	(2.26)	(1.30)	(1.77)	(-1.86)
TP x WI	0.03	0.25	0.26***	0.09	-0.01*
	(0.44)	(1.61)	(4.86)	(0.56)	(-1.72)
WF x WI	-0.09	-0.28**	0.02	0.00	-0.01**
	(-1.63)	(-2.24)	(0.42)	(.)	(-2.08)
R^2	0.320	0.940	0.937	0.762	0.063
Obs.	69,516	31,950	29,480	21,945	67,733
Baseline Controls	Yes	Yes	Yes	Yes	Yes
District FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Region× Year FE	✓	✓	✓	✓	✓

Note: p-values: *<0.1,**<0.05,***<0.01. The table shows the results of Equation 3.2. Only events are included where, in the subsample, at least 30 events have been recorded to avoid outlier events. All models are clustered at the province-year level. The outcomes are based on NDVI anomalies, log production of Cassava and Maize in tons, food insecurity and estimated GDP growth per capita by district. More context on the data generation is given in subsection C.7.

C.2 Additional Figures

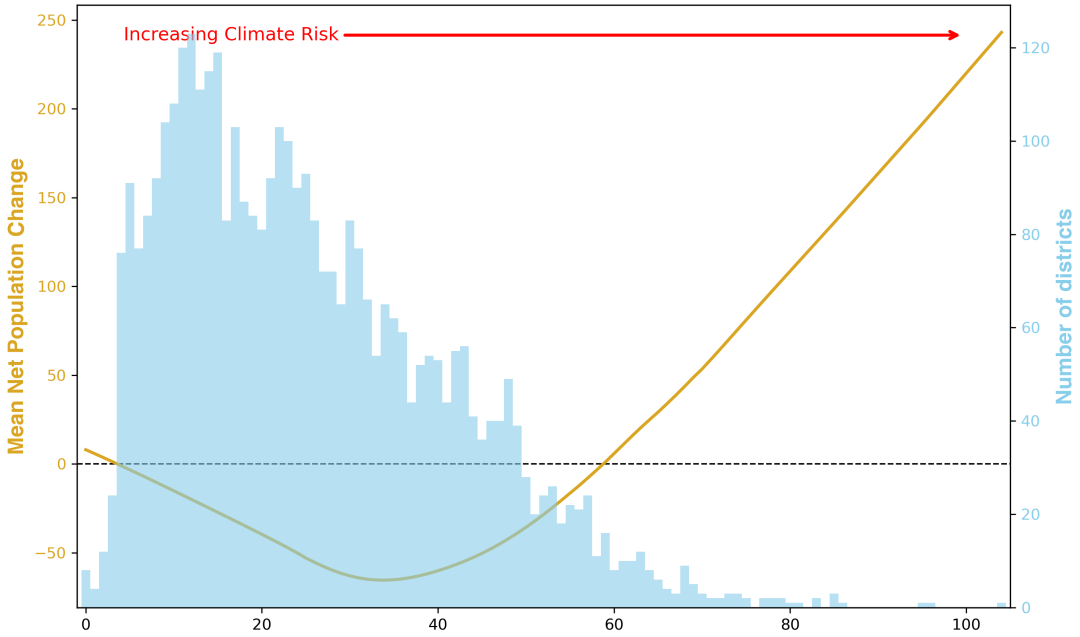

Fig. C.1. Net Population Change and Hazard Exposure

Note: The figure plots the net population change (in yellow on the left y-axis) against total exposure to hazard events (histogram plotted in blue). The line indicates the polynomial fitted trend line. Data on net population change is drawn from Niva et al. (2023) and hazard events are drawn from Reinhardt (2024a) and converted to binary sequence dummies using a threshold approach.

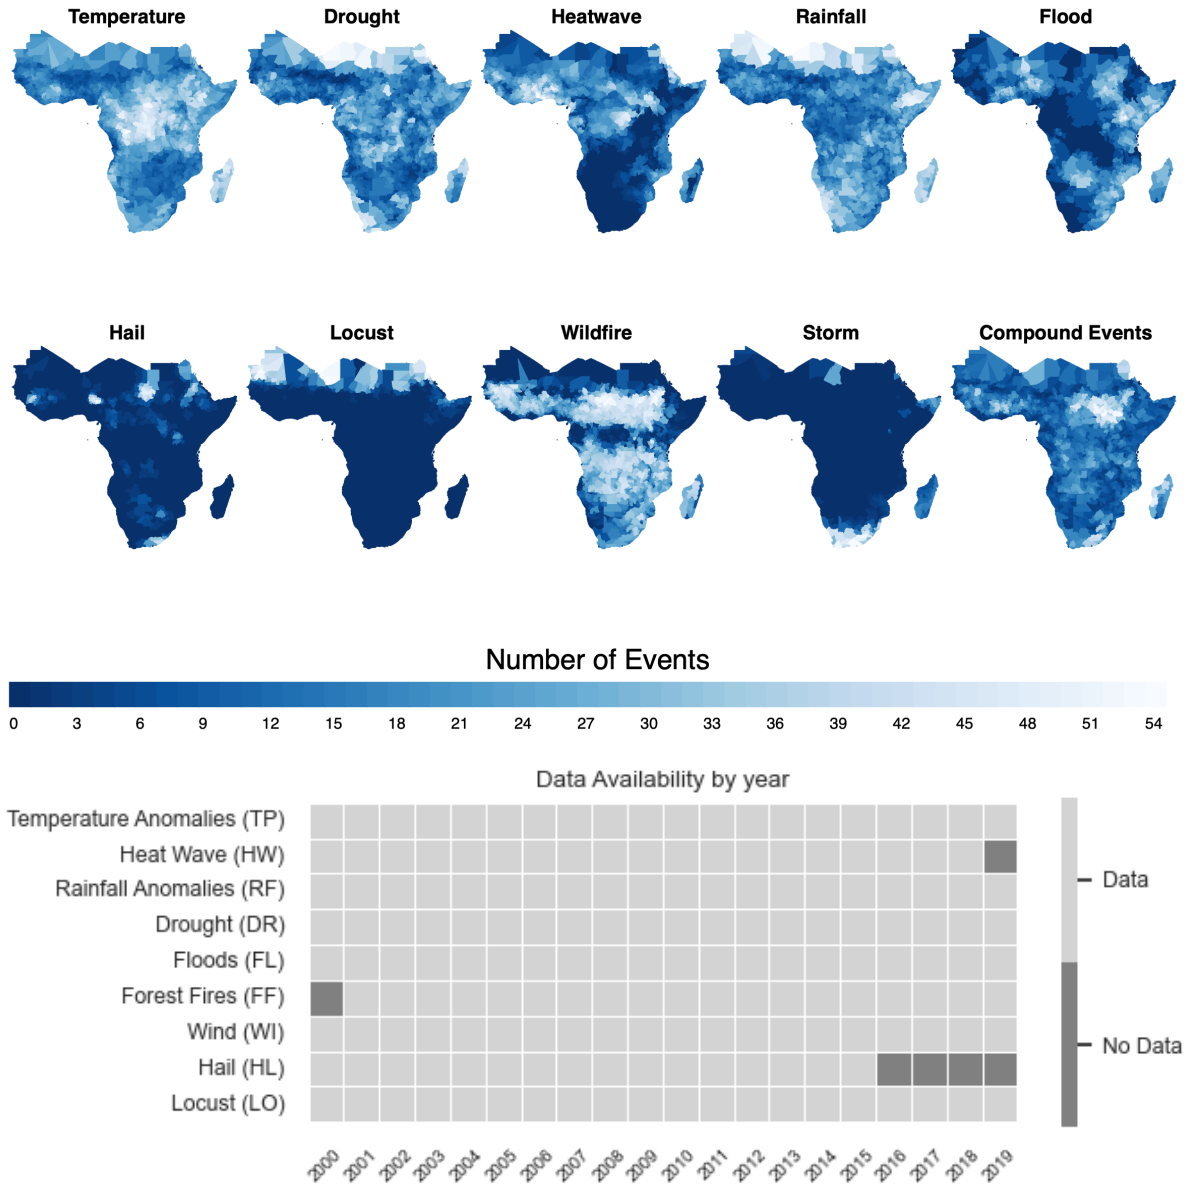

Fig. C.2. Spatio-Temporal Distribution of Shocks

Note: The upper figure plots for every hazard event the number of events recorded over 20 years by district. More white areas are strongly exposed. Compound events are defined as events where two different types of hazards are recorded within the same month. The lower figure plots the availability of input data in the dataset.

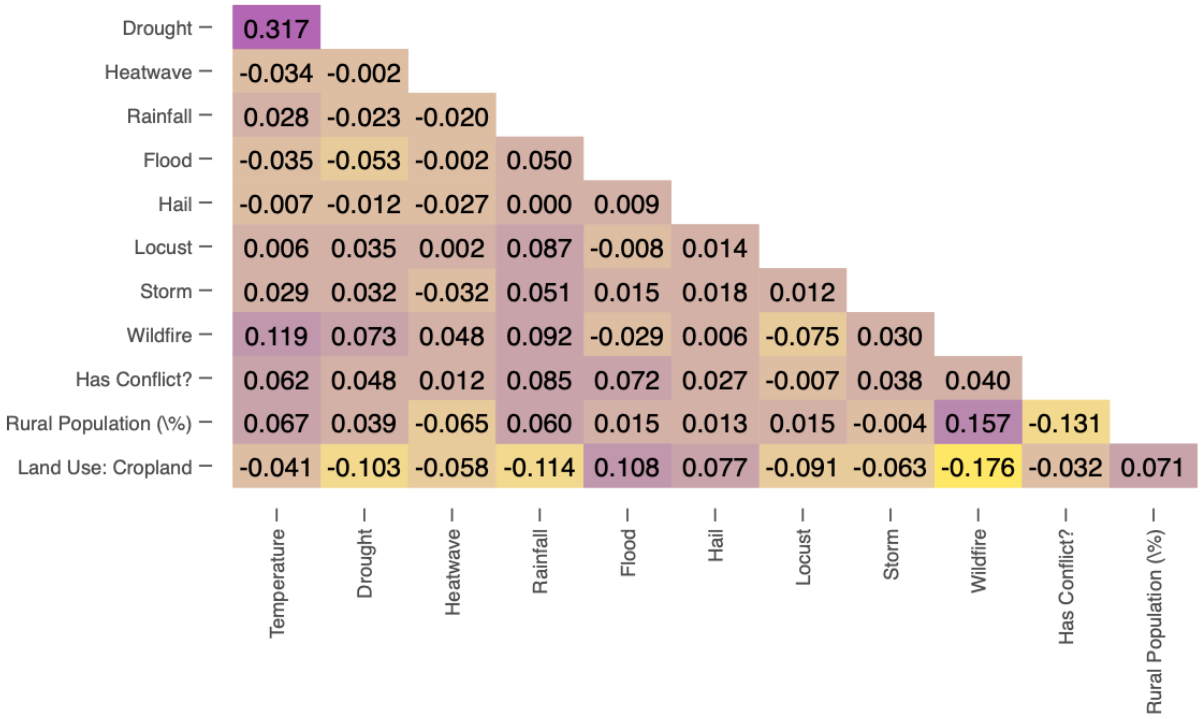

Fig. C.3. Correlation Matrix

Note: The matrix shows the standard correlation between individual coefficients and control variables.

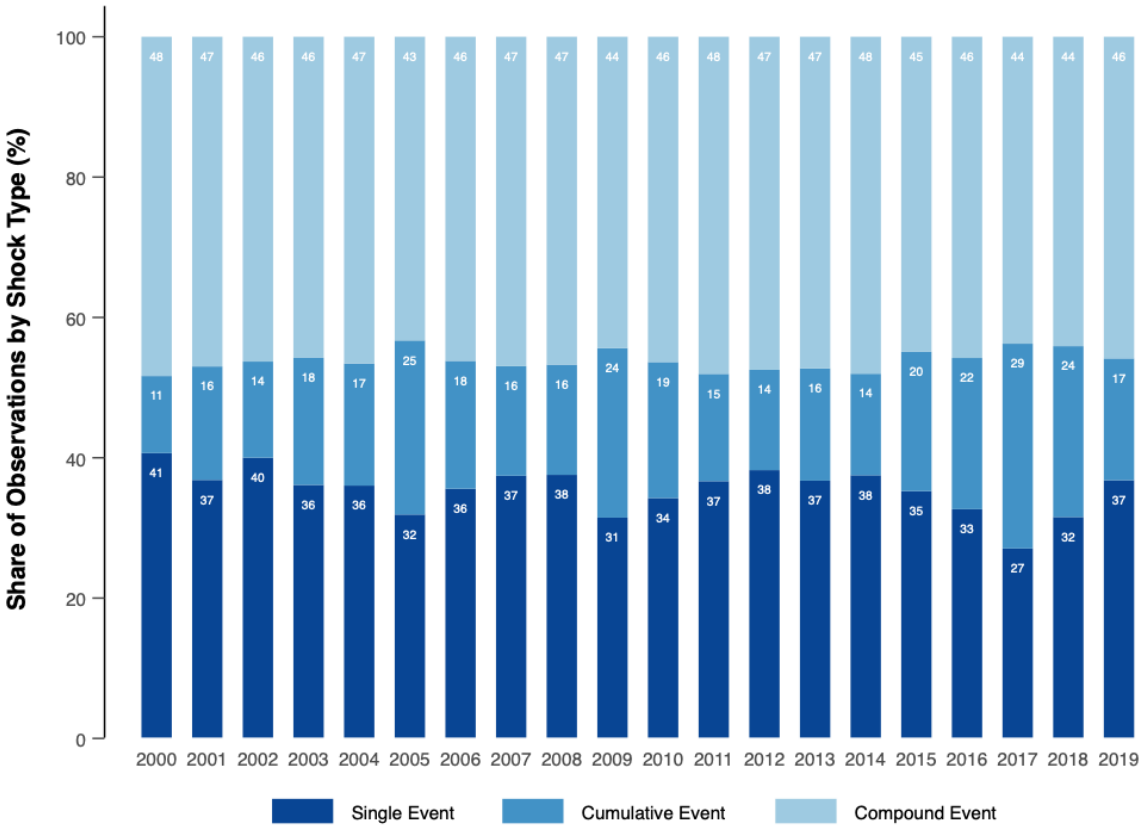

Fig. C.4. Fraction of Events By Type

Note: The figure shows of all recorded hazards by year, how common "compound" types as monthly interactions are. Hereby, the total number of events is divided into single events, having cumulative events (more than one event-sequence) and "compound" events.

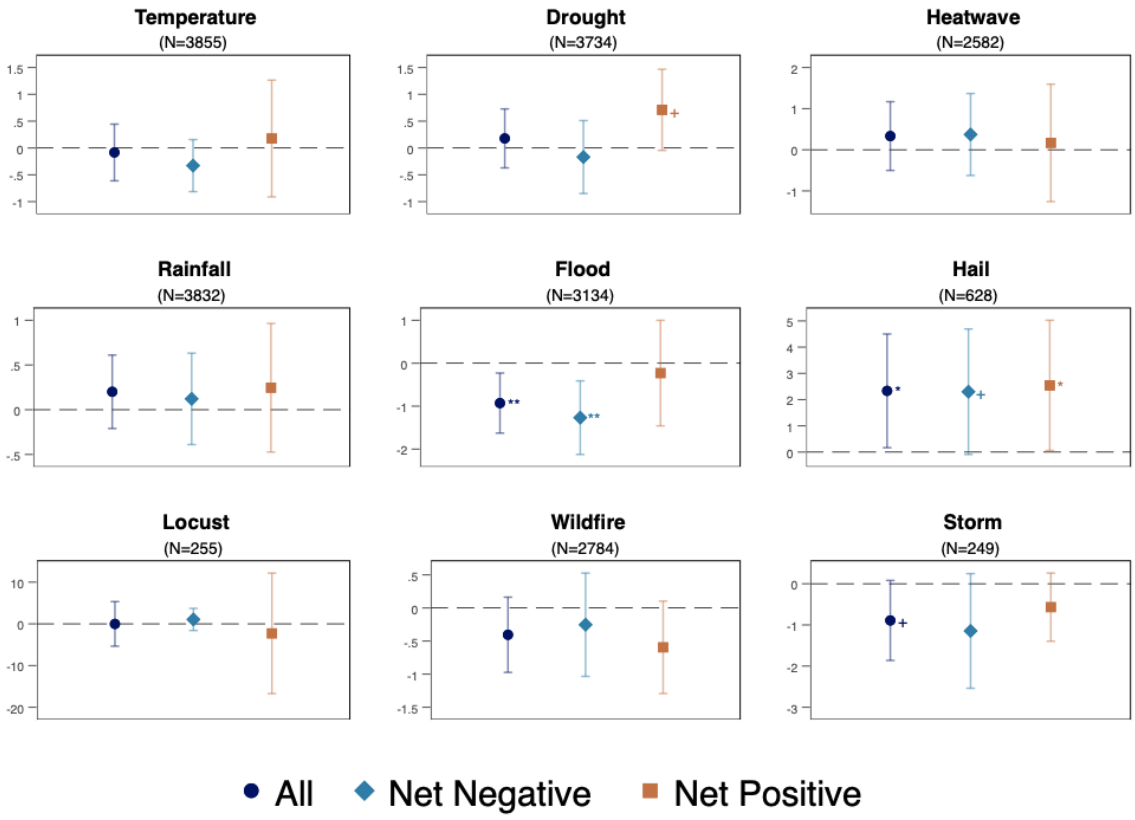

Fig. C.5. Baseline: Splitting into Net-Receivers and Emitters

Note: *<0.1, **<0.05 and ***<0.01 p-value. + p< 0.15. The figure shows the results splitting Equation 3.1 into net receivers (net positive) and net emitters (net negative), defined as the average over the covered 20-year period.

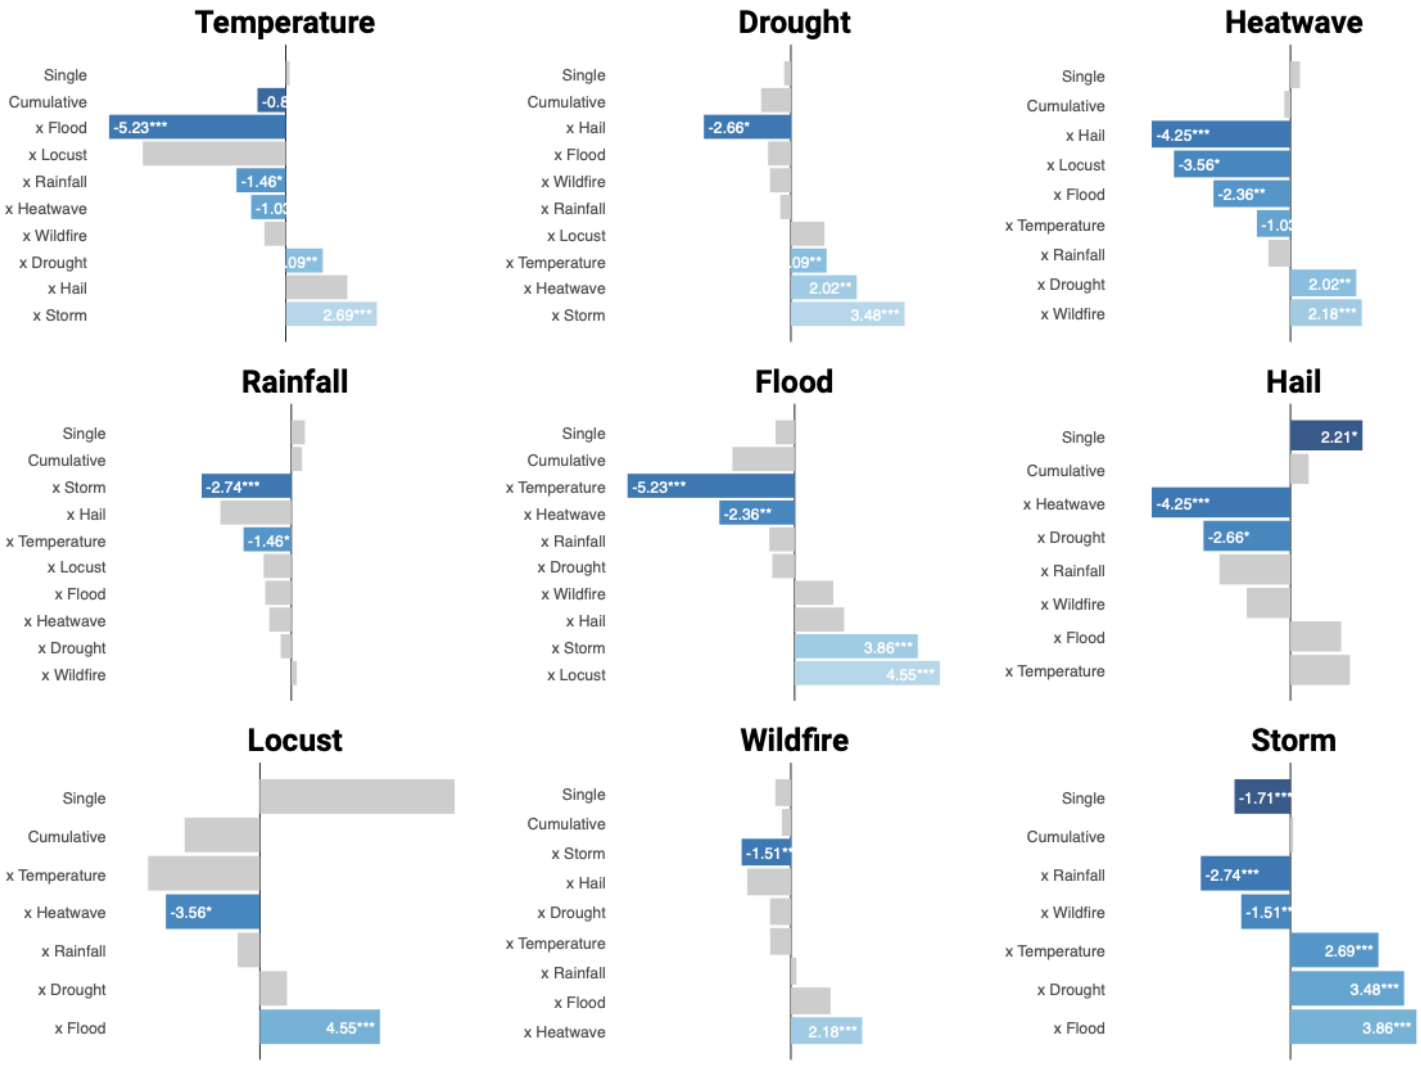

Fig. C.6. Effects by Hazard Type Full

Note: *<0.1, **<0.05 and ***<0.01 p-value.. The figure shows the marginal effects of being exposed to single events, cumulative events (more than one event of the same type within a year) or a compound event (both events occur in the same month). Insignificant results are shown as gray bars. All models include district-, year- and continental region-year fixed effects as well as controls for rural population, cropland share and presence of conflict. Standard errors are clustered at the province level.

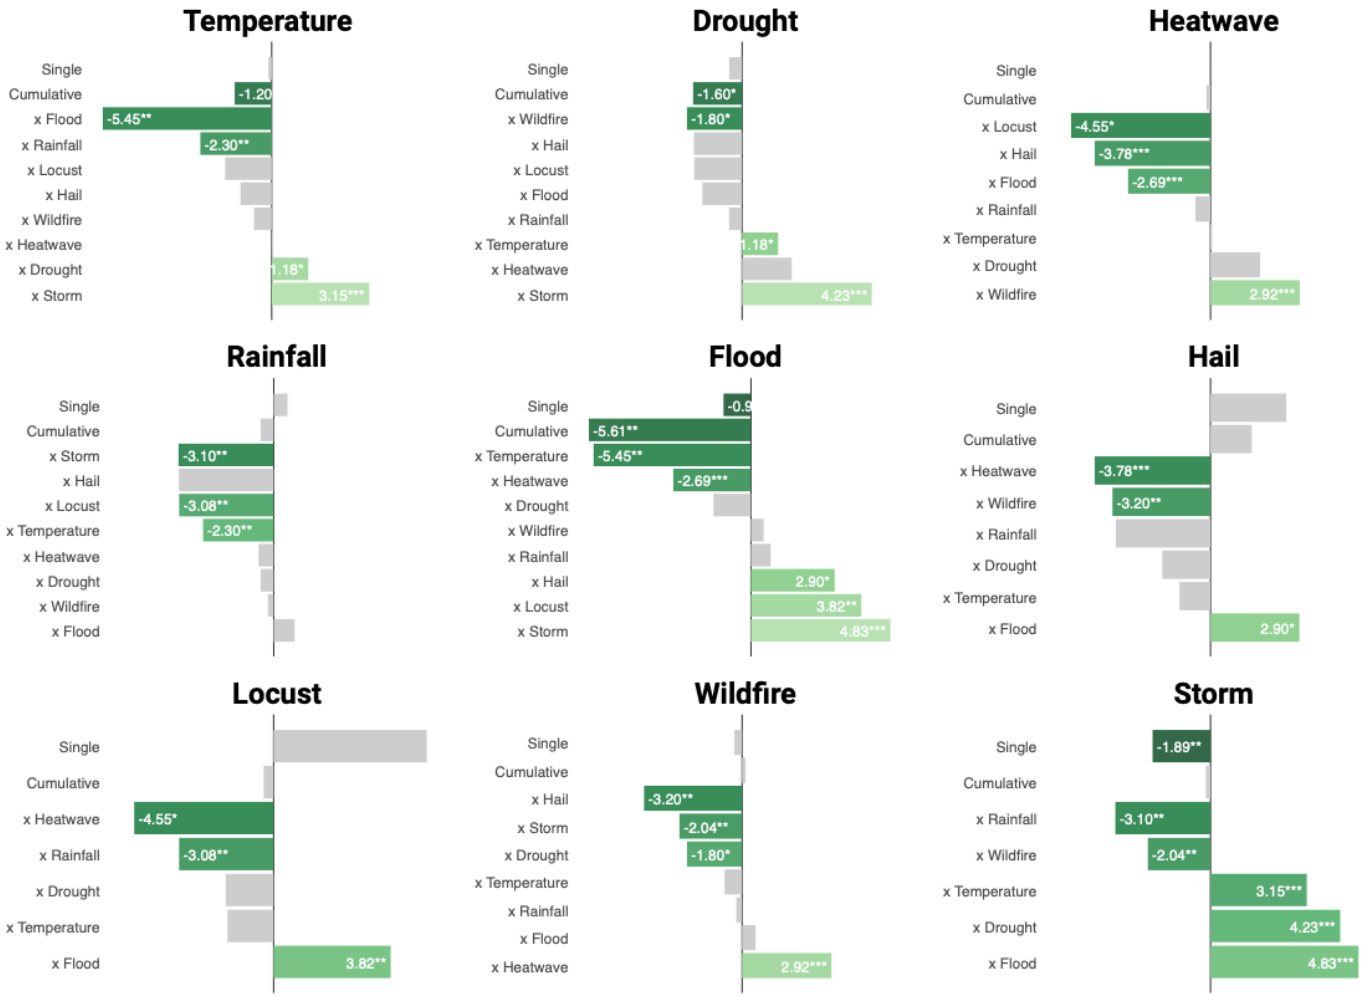

(a) Net Negative

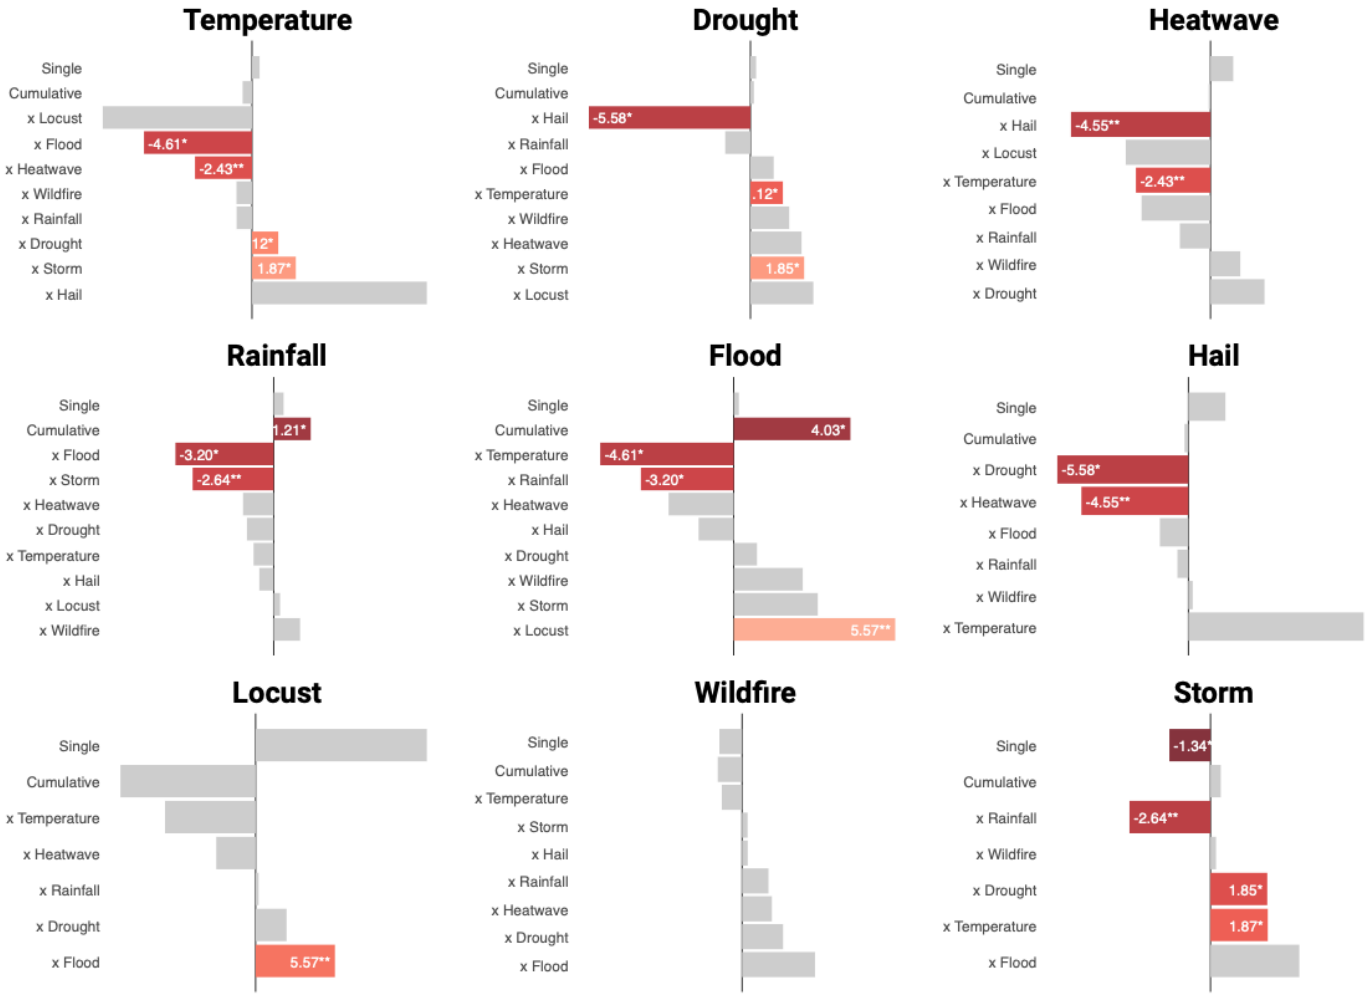

(b) Net Positive

Fig. C.7. Effects by Hazard Type and Emitting Type

Note: *<0.1, **<0.05 and ***<0.01 p-value. The figure displays marginal effects of single events, cumulative events (same-type events within a year), and compound events (both types in the same month). The sample is split by net population change over 20 years. Gray bars indicate insignificant results. Models include district, year, and region-year fixed effects, with controls for rural population, cropland share, and conflict presence. Standard errors are clustered at the province level.

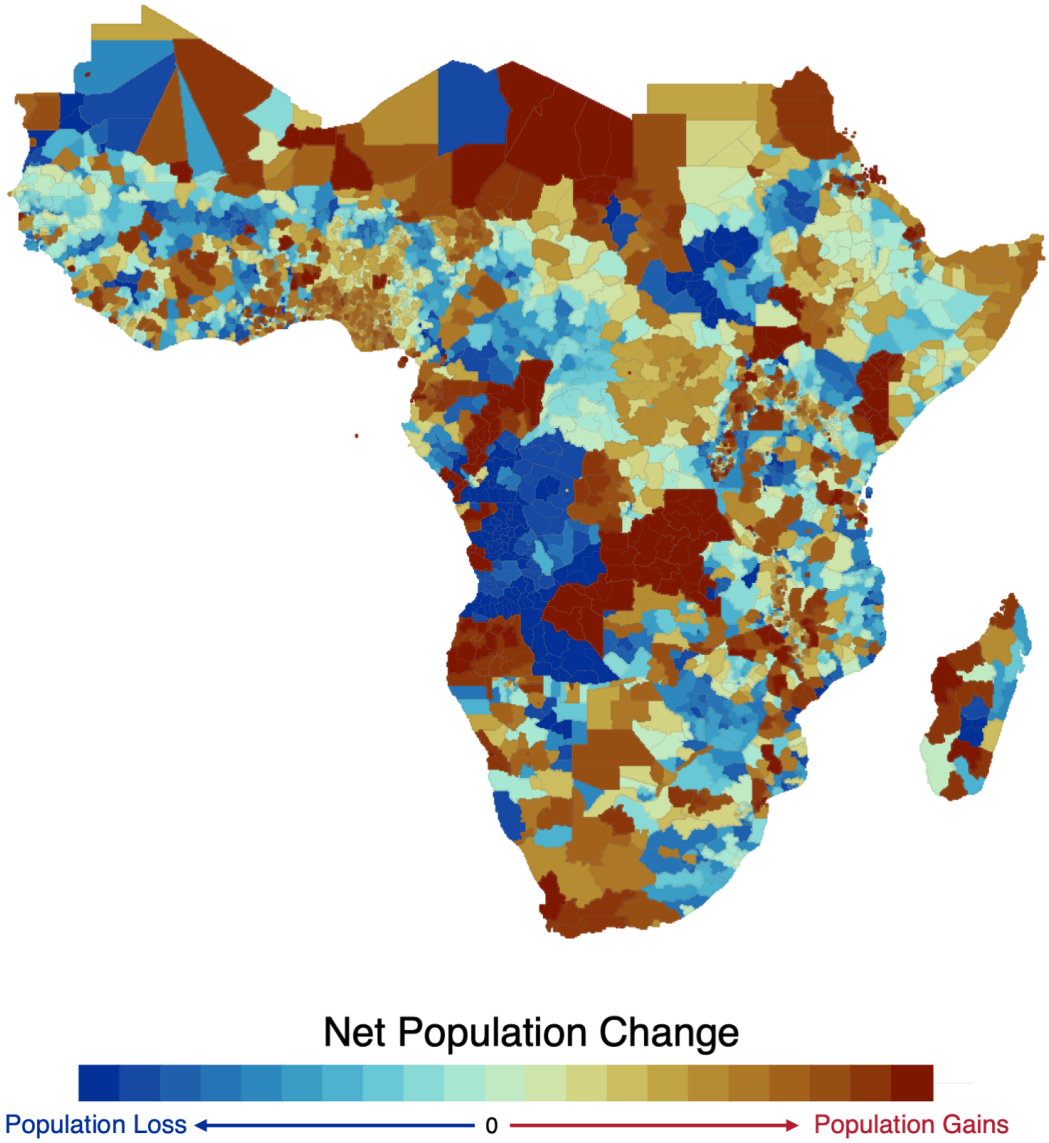

Fig. C.8. Net Emitting vs. Net Receiving Areas

Note: The map plots the 20-year average of net population change, which is used to classify net-emitting and net-receiving districts in Sub-Saharan Africa. The data is based on Niva et al. (2023).

C.3 Defining Compound Events

A significant part of this article categorizes the treatment into three types: single events, cumulative events, and compound events. In general, events are defined as mentioned in the Data Section. Events must exceed the predefined intensity threshold. The event itself is created as a unique sequence and accounts for the first month where a threshold is exceeded.

For compound events, a slightly modified approach is used. In simple terms, a compound event is defined as the first instance of simultaneous monthly overlap between two distinct types of hazards, where both events exceed their respective thresholds. In practice, this is accomplished through a five-step algorithm, which will be illustrated using an example of a flood-heatwave interaction in the Betsiboka district, Madagascar. First, an event sequence is created for each hazard. A sequence is initiated from the month in which the hazard threshold is exceeded and all consecutive months. In the example, the threshold for flooding is exceeded in February and remains in effect until April. Therefore, the period from February to April will be classified as a unique flood sequence. For a heatwave, the same is true, but starting in March and May. Each unique sequence run receives an identifier. In the example, the flood sequence receives the identifier (MDG-FL-30) and heatwave (MDG-HW-23). Next, we label for each pair the first month that overlaps. This shall avoid adding false positives, especially if sequences overlap for more than a single month. In the example, March will be the first overlapping month and thus identify a flood-heatwave interaction. The algorithm then eliminates individual overlapping runs and gives a unique identifier for the overlapping event. In the example, this eliminates MDG-FL-30 and MDG-HW-23, but adds a new dummy CP-FLHW-3 for the compound event. We eliminate both sequences as we want to avoid double-counting the same event once as a compound event and twice as a single event (for heatwave and flood). The remaining sequences that are not eliminated are finally counted and aggregated on an annual level. If, for example, there was another flood (MDG-FL-31) in October, the final aggregate will have one single flood event (the one in October) and one flood-heatwave interaction (the overlap in March between both events).

There are several shortcomings with this approach. First, we only identify 1-month overlaps. However, the effects of interactions might also be amplified if two events occur not necessarily in the same month, but in a 2- or 3-month window. This leaves room for improvement for future models. However, we include robustness interac-

tions of two types of hazards at the yearly level. This allows us, to some extent, to study how two hazards occurring within the year might impact outcomes. Second, using a one-month overlap does not allow for accounting for the order of the events. A flood following a drought might have different impacts than the reverse. This is because drought can cause the soil to harden, making areas less likely to absorb rain and potentially leading to stronger floods. This is a limitation that cannot be avoided in this setup. Lastly, we only account for bivariate combinations. Having already used bivariate combinations, we have more than 30 compound event types. If we were to allow for all possible unique combinations, our final model would need to account for 511 compound event types, which would likely result in overfitting our regression model. Future research should identify the most critical combinations based on discussions with climate scientists and agricultural researchers to account for the most significant interactive types.

C.4 Robustness Checks of the Baseline

Testing for Functional Form of Baseline To test for the presence of nonlinearity, we assess if the functional form of the baseline model is correctly specified. Table Table C.5 shows the use of monthly counts rather than sequences as treatment variables. Panel A adds duration in linear terms, while Panel B adds quadratic terms. It appears that the number of months with temperature anomalies has a negative linear effect on the net population change. Flood and hail show a similar direction as in the baseline, albeit not significant. Once quadratic terms are added they turn significant again. Hereby, more prolonged floods turn coefficients more positive, while the opposite is true for hail. For flood, the turning point would be around 1.7 months of exposure. This exposure is reached for 25% of all flood events; therefore, for a quarter, our model averages and potentially masks the true effects of flood events. For hail, the computed turning point is 1.6 months, which is only reached in around 7% of all events. In both cases, most of the events show the described baseline results. Yet, we acknowledge that in around a quarter of the events, we might find a potentially misspecified structural form, where a nonlinear structure might have add value.

Testing for Spatial Spillovers

To test for spatial spillovers, we rely on a Spatial Durbin Error Model following Anselin (1988). In contrast to the standard two-way fixed effects model, SDEMs take the spatial relationship between districts into account. To capture spatial spillovers commonly

Table C.5
Robustness: Changing Treatment Definition

	Temperature	Drought	Heatwave	Rainfall	Flood	Hail	Locust	Wildfire	Storm
Panel A: No Squared Terms									
Duration (Months) (t-1)	-0.26** [-0.47,-0.05]	0.07 [-0.16,0.30]	-0.01 [-0.27,0.25]	0.10 [-0.20,0.39]	-0.03 [-0.39,0.33]	1.05 [-0.01,2.10]	1.23 [-1.80,4.26]	-0.25 [-0.52,0.02]	-0.03 [-0.20,0.14]
Panel B: Squared Terms									
Duration (Months) (t-1)	-0.24 [-0.69,0.22]	0.11 [-0.49,0.72]	0.29 [-0.14,0.73]	0.39 [-0.31,1.08]	-1.22*** [-1.98,-0.46]	3.28** [0.65,5.90]	-7.36 [-20.92,6.19]	-0.17 [-0.71,0.37]	-0.24 [-1.15,0.66]
Duration (Months) sq. (t-1)	-0.01 [-0.10,0.09]	-0.01 [-0.17,0.15]	-0.06* [-0.10,-0.01]	-0.16 [-0.50,0.18]	0.37** [0.10,0.64]	-1.02** [-1.75,-0.30]	1.93 [-1.53,5.39]	-0.00 [-0.15,0.14]	0.03 [-0.16,0.23]
Obs.	69,516	69,516	69,516	69,516	69,516	69,516	69,516	69,516	69,516
Districts	3855	3734	2582	3832	3134	628	255	2784	249
Unit-FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year-FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Region-Year-FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls?	✓	✓	✓	✓	✓	✓	✓	✓	✓
Control for Hazards?	✓	✓	✓	✓	✓	✓	✓	✓	✓

Note: *<0.1, **<0.05 and ***<0.01 p-value. All models include district, year and region-year fixed effects. As controls, lagged conflict count, lagged rural population share, and lagged share of cropland are added.

two spatial econometric models are found: Spatial Durbin Models (SDMs) or Spatial Durbin Error Models (SDEMs). The former allows for feedback effects of the outcome by making the spatially lagged dependent variable a core part of the regression model. In contrast, we opt for a SDEM model that assumes that net population change predominantly responds to spatially correlated shocks, but not directly to neighboring districts' net population change trends. SDEMs build spatial lags for treatments and covariates and also account for latent spillovers by incorporating spatial autocorrelation into the standard errors.

The SDEM model is highly dependent on the setting of a spatial weighting matrix, which represents the influence of neighbors on each other. There are many ways in which neighbors can be weighted. We opt for a row-standardized Queen contiguity spatial weighting matrix, which is common in this literature. This matrix is now called W. Then the SDEM model of the following form is run:

$$Y_{ijcrt} = \beta_0 + \beta_1 D^{(k)}_{ijcr,t-1} + \gamma X_{ijcr,t-2} + \theta_1 W D^{(k)}_{ijcr,t-1} + \delta W X_{ijcr,t-2} + \lambda_i + \lambda_t + u_{ijcrt} \tag{3.4}$$

$$u_{ijcrt} = \rho \sum_{j'} w_{jj'} u_{ij'crt} + \xi_{ijcrt} \tag{3.5}$$

The implementation is based on the `spreg` Python package (Amaral et al., 2025) and follows a form similar to the models of Felbermayr, Gröschl, et al. (2022). The difference to the baseline model is that here no background stressors (T) are added. This is done to reduce noise in the regression model. Additionally, region-year fixed effects are excluded because they may absorb part of the variation, especially for districts located along the region's boundaries. The $D^{(k)}_{ijcr,t-1}$ measures the direct effect of a hazard event of type k on the outcome, while the $WD^{(k)}_{ijcr,t-1}$ measures the spillover effect for the nearby districts. The updated error term u_{ijcrt} controls the spatial autocorrelation in the error term and shall account for omitted unobservables spatially. To make the results more comparable to the baseline, we re-run a modified baseline regression that is identical to the above equation, with the exception that WD and WX are omitted, and standard errors are clustered at a district level.

The results are shown in Table C.6. Similarly to the baseline in Equation 3.1, floods and hail are significant. However, because of changes in the fixed effects, wildfires and storms also become significant. More importantly is Panel B. It appears that temperature anomalies and droughts have significant spillover effects, where the effects

in neighboring districts offset the reduction in net population change. The effects of floods remain significant and negative, with no observed spillover effects. The finding that some of the results of the two-way fixed-effects model might be biased due to hidden spatial spillovers justifies why, in the baseline, a province clustering rather than a district clustering is used. However, some residual spillovers may still exist. Wildfires primarily have local effects that are not influenced by spillovers, highlighting that the reallocation of the local population due to nearby wildfires is unlikely. For locust events, both direct and spillover effects are large and positive, in contrast to the insignificant TWFE estimates. This reflects spatial confounding in the baseline model. The strong effects are likely due to the high spatial concentration of desert locust swarms around a small part of Africa. The baseline TWFE averages over the whole region, whereas the SDEM accounts for the strong effects found within the small region that is concerned by locust swarms. For droughts, if the districts themselves are exposed, no significant impact on net population change is recorded. Yet, if neighboring districts are treated, net population increases indicate some form of migration or displacement. Overall, these findings reveal that, for some, albeit not all, effects, spatial spillovers exist and should be taken into account.

Temporal Dynamics The baseline model assumes that treatment takes precisely one period to show effects. Effects, however, might be dynamic and carry over multiple time periods. Also, anticipation and pre-trends might be present and potentially influence the coefficients. To test if our model is driven by treatment or has dynamic carryover effects, we estimate an event-study version of our baseline model relying on the DiD_l estimator of De Chaisemartin and d'Haultfoeuille (2024). To estimate treatment effects that allow for treatment reversal, the parameters of interest are

$$ATT(d,t) = \mathbb{E}[Y_t(d) - Y_t(0) \mid G = d]$$

where the average treatment effect of the treated is split into treatment sequences (d) and periods (t). Following Baker et al. (2025) the estimator "staggerizes" such treatment sequences (d) following the first time the districts start to be treated. The DiD estimate is normalized by the number of treated periods. In our setup, we use the same baseline controls as in Equation 3.1, but we modify the treatment definition slightly. In contrast to the baseline, where events need to affect at least 20% of a district, here we set a higher threshold of more than 80%. This is introduced because otherwise our sample of "clean" pretreatment periods becomes so small that estimation is not feasible without contamination of previous districts. Furthermore, this approach reduces potential

Table C.6
Robustness: Spatial Durbin Error Model

	Temperature	Drought	Heatwave	Rainfall	Flood	Hail	Locust	Wildfire	Storm
Panel A: Modified Baseline Model									
Has Treat (t-1)?	-0.28	-0.18	0.08	0.18	-0.71**	1.84**	-0.45	-0.89***	-0.70**
	(0.26)	(0.25)	(0.31)	(0.22)	(0.29)	(0.78)	(3.19)	(0.29)	(0.32)
Panel B: SDEM Model									
Has Treat (t-1)? t	-0.45*	-0.58*	0.31	0.40	-0.77***	0.75	21.76***	-1.57***	1.42
	(0.24)	(0.35)	(0.35)	(0.34)	(0.38)	(0.88)	(5.15)	(0.39)	(1.51)
W × Has Treat (t-1)?	1.13***	1.30**	0.83	1.20	1.40	-0.90	46.20***	-1.21	5.04
	(0.50)	(0.78)	(0.70)	(0.79)	(0.87)	(2.34)	(16.27)	(1.04)	(4.25)
Obs.	69,534	69,534	69,534	69,534	69,534	69,534	69,534	69,534	69,534
Unit-FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year-FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls*?	✓	✓	✓	✓	✓	✓	✓	✓	✓

Note: *<0.1, **<0.05 and ***<0.01 p-value. Panel A shows a modified baseline regression with a similar setup to the SDEM model. Panel B implements an SDEM model relying on Amaral et al. (2025).

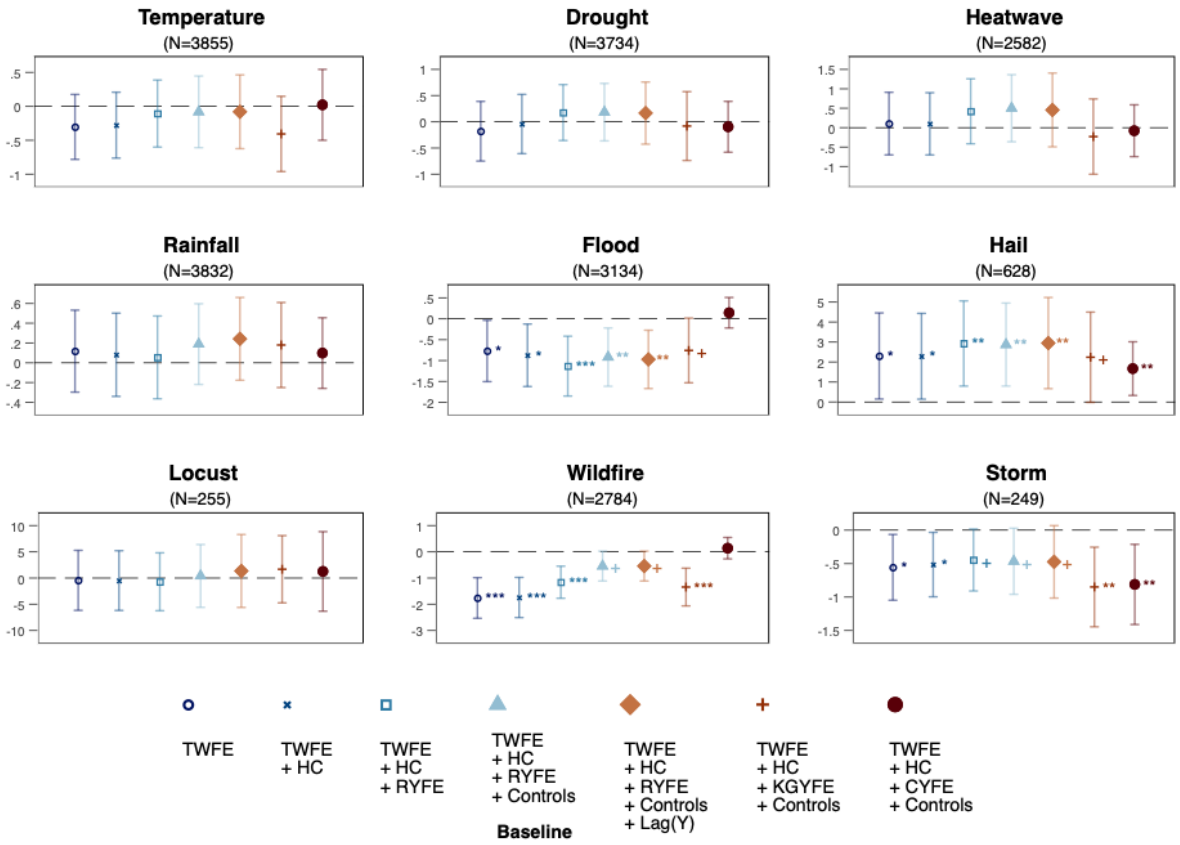

Fig. C.9. Robustness: Selection into Treatment

Note: *<0.1, **<0.05 and ***<0.01 p-value. The model presents the results of the baseline specifications, with the difference that each regression is conditioned only on districts that experience at least one event of the hazard type over 20 years.

missclassification or events that might have been too weak to take effect. Moreover, we exclude the set of background stressors, as in some treated groups, the first-difference regression of the outcome on control variables and fixed effects, which is part of the estimation, cannot be estimated due to insufficient variation. Here, the number of observations is smaller than the number of control variables. To avoid errors and bias in the estimates, we thus exclude the background stressors from the model.

The results are shown in Figure C.10. Locust swarms, drought, temperature and hail show no significant effects or clear pre-trends. The heat wave seems to increase the net population change from two years after treatment onward. Rainfall, in contrast to the baseline, has a strong and significant effect. We suspect that this extreme effect is likely driven by extreme outlier events, given that rainfall anomalies might be more local. Further investigation shows that rainfall events extending for more than

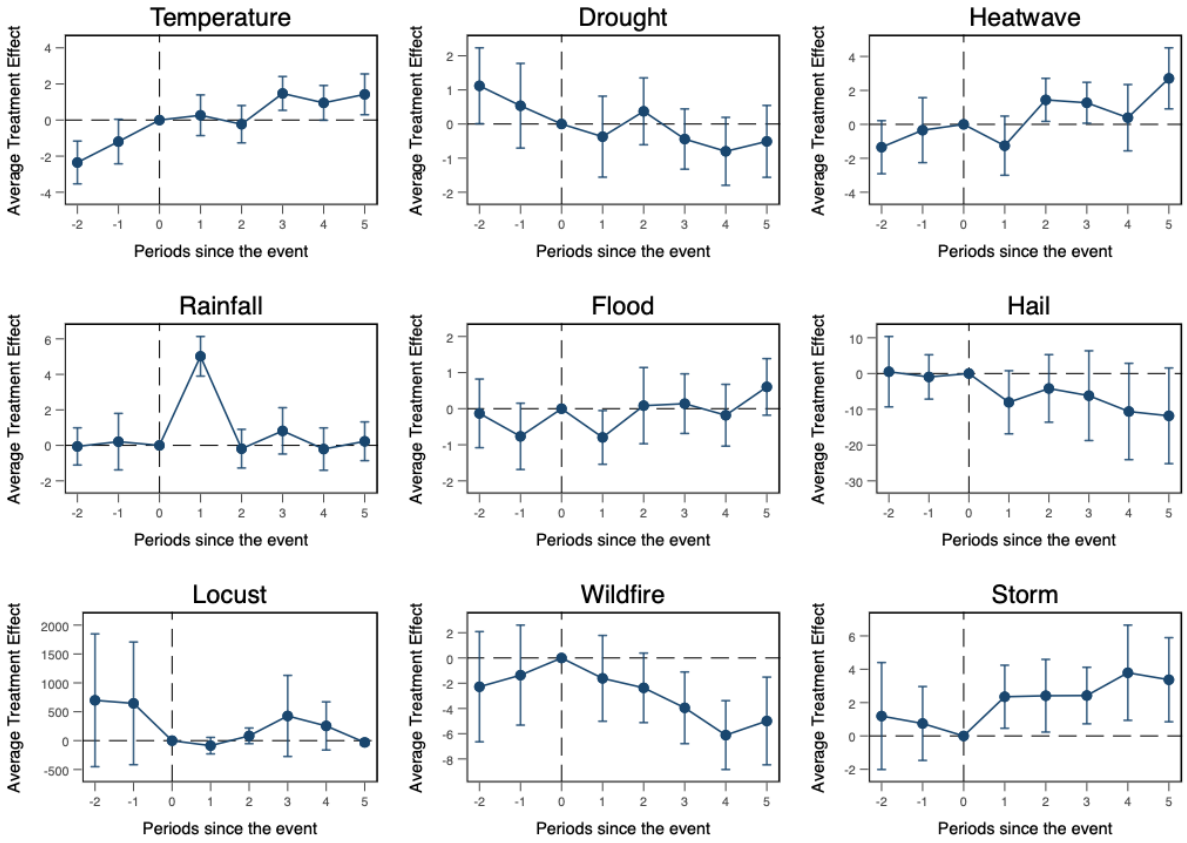

Fig. C.10. Robustness: Dynamic Treatment Effects

Note: *<0.1, **<0.05 and ***<0.01 p-value.. The shown estimates present an event study using individual treatment on a district level and estimate the DiD_l estimator as described by De Chaisemartin and d'Haultfoeuille (2024). Standard errors are clustered at the district level, following the recommendations of De Chaisemartin and d'Haultfoeuille (2020). Confidence intervals reflect 90% intervals.

80% of a district represent only around 17% of all recorded rainfall events. Thus, the strong response is likely driven by severe events in the upper tail. Flood shows negative effects in the year after a flood, consistent with the baseline. However, the effects are only marginally below the 10% threshold. The clearest patterns are observed for wildfires and storms, where storms increase the net population and wildfires reduce it. Plausible mechanisms might include the fact that wildfire destroys infrastructure and causes income losses, thereby driving people out of districts. However, if human-caused, it could also be indicative of slash-and-burn practices that result in larger crop yields and thereby enable households to send migrants. Storms, in contrast, may reduce out-migration and thereby increase net population change in areas with negative

net population change, especially if storms cause households to become trapped in poverty. Overall, the results show that flood is also, albeit with little significance, negative in the event study; however, our baseline model may underestimate the dynamic treatment effects, especially for wildfires and storms. Lastly, at least for the chosen two periods apart from temperature, no pretrends are shown, potentially indicating that anticipation plays a minor role. These results should be interpreted with caution, especially since the estimator by De Chaisemartin and d'Haultfoeuille (2024) relies, like our model, on the Stable Treatment Unit Value assumption, which we have previously documented as potentially being violated.

Microlevel Validation The analysis of the baseline is conducted on a fairly aggregate level. Following Hoffmann, Šedová, and Vinke (2021), many studies on the migratory response to shocks are conducted at the country macro-level or the household micro-level. This article employs an intermediate level of aggregation at the district-unit level. This might capture broad population trends but does not necessarily capture the household response to a shock. Although not at the center of this study, we provide suggestive evidence on the effects at the household level by reproducing a model similar to that of Di Falco et al. (2024). Their article focuses on the effects of cumulative droughts on household decisions to migrate. Instead of using all Living Standards Measurement Survey Integrated Surveys on Agriculture (LSMS-ISA) of the World Bank, we restrict this robustness test solely on Nigeria (National Bureau of Statistics (NBS), 2014), given its importance of population and exposure to many of our hazards. To create the data, the geolocations of the LSMS-ISA clusters are intersected with the grid taken from Reinhardt (2024a), and then a panel Logit model of the following form is estimated:

$$\log\left(\frac{\Pr(M_{iht}=1)}{1-\Pr(M_{iht}=1)}\right) = \alpha + \sum_{k=1}^{K} \beta_k D^{(k)}_{ih,t-1} + \gamma X_{iht} + \phi_h + \epsilon_{ht} \tag{3.6}$$

The outcome M_{it} approximates the likelihood that an individual leaves the household without returning[13]. To avoid death of household members, we retain a sample of only households that do not lose any member over the three waves covered. Thus, we avoid capturing persons who might have died as a result of the direct or indirect effects of an extreme event. The $\sum_k^K \beta_k D^{(k)}_{ih,t-1}$ test the marginal effects on the likelihood of having a migrant if in the previous year a shock of this kind occurred. All are added at once

[13] Following the question: "Is person X still a member of the household?"

to account for the potential presence of other hazards. In additional specifications, we replace the dummy with an alternative dummy that indicates whether there was cumulative exposure in previous years (e.g. if floods occurred in both prior years). Only six of the nine hazards are included since some barely occur in Nigeria. The vector $X_{i,t}$ controls the same household and individual-level determinants than in Di Falco et al. (2024) (amongst others gender, household size, or non-agricultural income). In contrast to Di Falco et al. (2024) we estimate a logit regression and exclude country-level estimates on natural and non-natural disaster exposure given that we only compare effects within the same rather than between different countries. The results can be seen in Figure C.11. On a household level, having had a temperature anomaly in the previous year (red) does not change the odds of migration. Drought and flood significantly reduce the odds of having a family member who is a migrant. Precipitation alone, in contrast, increases the odds, making a household 1.2 times more likely to have a migrant. A likely explanation is the reduction in household income after droughts and floods, which are known to affect agriculture and physical assets. Rainfall, on the other hand, may increase crop yields, thereby contributing to higher income. In contrast to the baseline, no significant effects are found for hail, likely because hail is not extremely widespread in Nigeria. Overall, the household results reveal that for one country tested, flood and drought tend to reduce decisions to migrate. This does not directly contradict the baseline results. First, it is unclear where do people migrate to. Depending on whether they stay within the district or move farther away, this will influence the net population change. If they move further away, droughts and floods will reduce net population change in net-emitting regions and increase it in net-receiving areas. We do not find this in our baseline. However, the baseline also covers the whole subcontinent and although such effects might not be found in Nigeria, there might be such effects in other parts of Africa.

The other two coefficients (yellow and blue) in Figure C.11 test whether cumulative effects matter for households' decisions. In other words, it is tested if cumulative exposure over multiple years (e.g. having two consecutive years with severe droughts) influences the household decision. The effects of floods and droughts continue to reduce the likelihood of migration. Now, temperature anomalies reduce the likelihood of migration, while repeated forest fires increase the odds of having a migrant in the household. Potentially, a single bad year might be captured, but repeated exposure to the same event causes effects to vanish. In contrast to Di Falco et al. (2024), the propensity to migrate decreases with drought exposure, but this is likely due to the different

definition of treatment, the different model, and the set of controls. The results of cumulative exposure highlight that another channel, which is not at the center of this study, matters. This type of "compound" event is examined in Di Falco et al. (2024) and Petrova et al. (2025) and is not central to our study. For completeness, we include temporal compounding events to the baseline in Figure C.12 that capture having had a type of event over consecutive years. Cumulative drought, floods, and storms seem to matter, especially for drought and flood in net-emitting areas. Hail only matters in the year of treatment. Overall, this reveals that indeed, temporal cumulation of events tends to also matter, in line with findings from the microlevel test. This extension is descriptive; however, we would like to acknowledge that cumulative, repeated events over time may have a significant influence on household decisions regarding the presence of a permanent migrant and, consequently, on net population changes.

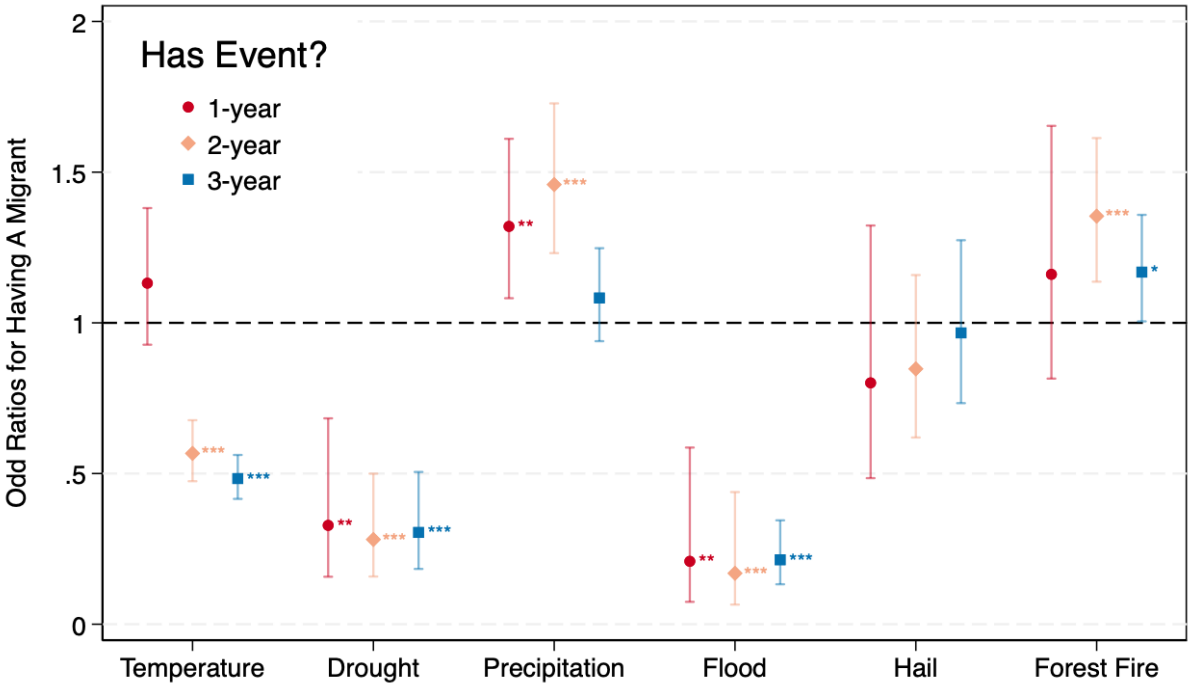

Fig. C.11. Robustness: Validation with LSMS

Note: *<0.1, **<0.05 and ***<0.01 p-value. The model estimates the model described in the paragraphs above. The data is based on LSMS-ISA waves 1 to 3 of Nigeria. The model adds all hazards jointly and controls for household and individual statistics. Regressions include household fixed effects and household-clustered standard errors. The different events capture having an event of type k in the previous year, over 2 years, and 3 years.

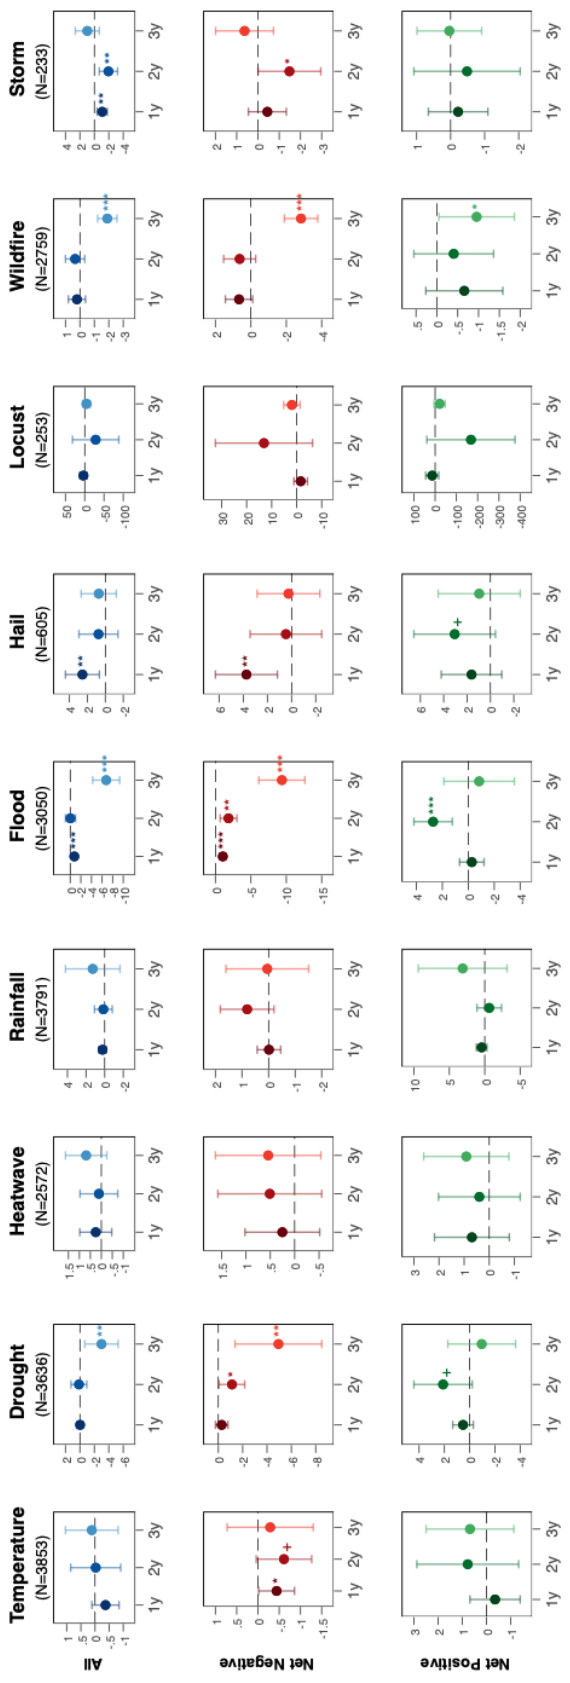

Fig. C.12. Robustness: Cumulative Events by Hazard Type

Note: *<0.1, **<0.05 and ***<0.01 p-value. The model estimates a similar model to Equation 3.1 with the exception that dummies $D^{(k)}$ are now categorical capturing "having event", "having event in 2 consecutive years" and "having event in 3 consecutive years". Controls, fixed effects and standard error clustering are the same. The upper panel uses the full sample, the lower two panels split the data into districts that have over a 20-year period net-negative and net-positive population change.

Other Robustness Tests To further validate the results, three additional robustness checks are included. First, Figure C.13 checks if results are driven by outliers of districts in the dependent variable. Thus, the net population change variable is winsorized using winsorization at 1% and 5%. The results do not change substantially in significance or direction.

Next, Figure C.14 investigates another change in the definition of treatment. Before, it was assessed whether our definition of upper-tail events might mask some underlying nonlinearities. A second condition that must be met is that a certain fraction of a district must be affected. This threshold is arbitrarily set to 20% of the physical area. This can be problematic if the area sizes of districts vary: 20% in a small district corresponds to a much smaller affected area than in large districts. This could include a bias if the treatment therefore captures different levels of exposure to events. To assess this, Figure C.14 varies the area threshold over several dimensions. In the upper panel, for each hazard, the thresholds ranging from marginally above 0% to 100% of the affected physical area are shown. The results using the chosen threshold of 20% are generally robust to changes in the cut-off. For hail, the most significant cut-off is around the chosen 20% level. Larger thresholds for hail result in a larger coefficient, but one that is less significant, below a 10% level, likely due to a smaller treatment sample. For storms, especially those with a 40% threshold, the coefficient becomes marginally significant at the 10% level, whereas all others fall below the traditional significance levels. In particular, locust swarms appear to have a significant coefficient beyond 80% of the affected district. However, only six districts have this intense exposure, making it impossible to draw a valid inference. Rainfall anomalies have a significant effect if more than 60% of a district is exposed. To provide further validation of the threshold 20%, we apply alternative weighting schemes: the share of population of a district affected and the relative cropland share affected. The results are comparable to those obtained using area weighting. With population and cropland shares, flood remains significant, whereas hail falls marginally below the 10% level. The same critique can be applied using population or cropland weighting, namely that 20% in districts with higher population density might show larger effects than in areas with lower population density. However, we believe that the distribution of population and cropland is not uniform within a given district, in contrast to its physical area. Thus, given that the results are comparable, it adds confidence that the effects are indeed relevant. Furthermore, as our baseline model includes district-fixed effects, it accounts for the total area of a district, thereby mitigating concerns that size itself is the primary driver of the treat-

ment effects. Overall, the results indicate that the main findings are not fundamentally driven by the chosen 20% threshold in the baseline specification.

Lastly, the data might be noisier for district-level data than for province-level estimates. Figure C.15 aggregates treatment and results at the provincial level and also applies a 20% threshold for affected land. Floods remain significant, especially for net-emitting provinces, but hail and storms lose their importance. In contrast, wildfires show similar effects to floods, again in net-emitting regions. Given that 20% of a province has substantially more burned area than 20% of a district, this strong effect might be driven by some extreme wildfire events that induce displacement. Finally, if large parts of provinces are affected by droughts, net receiving districts (with a positive net population change) receive a larger inflow of population. This aligns with the findings of the literature, which suggest that droughts may cause displacement in urban areas, which, on average, are net receiving areas.

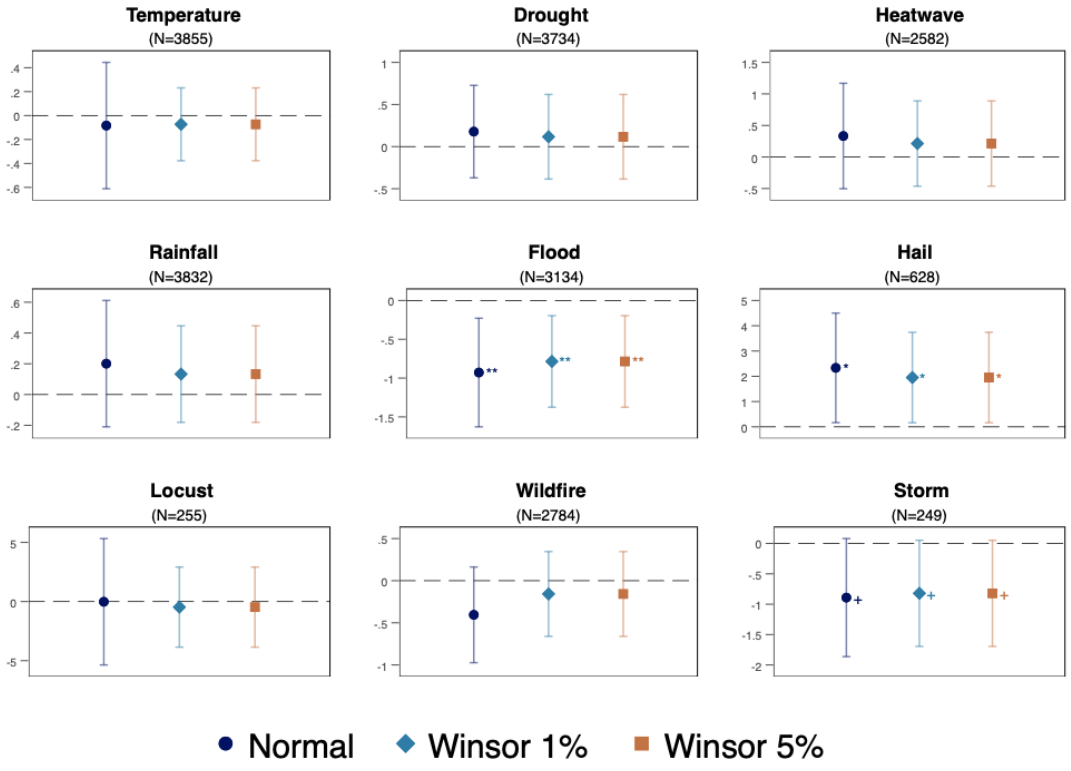

Fig. C.13. Robustness: Sample Dependence

Note: *<0.1, **<0.05 and ***<0.01 p-value. The model estimates Equation 3.1 but winsor the dependent variable by 1% or 5%.

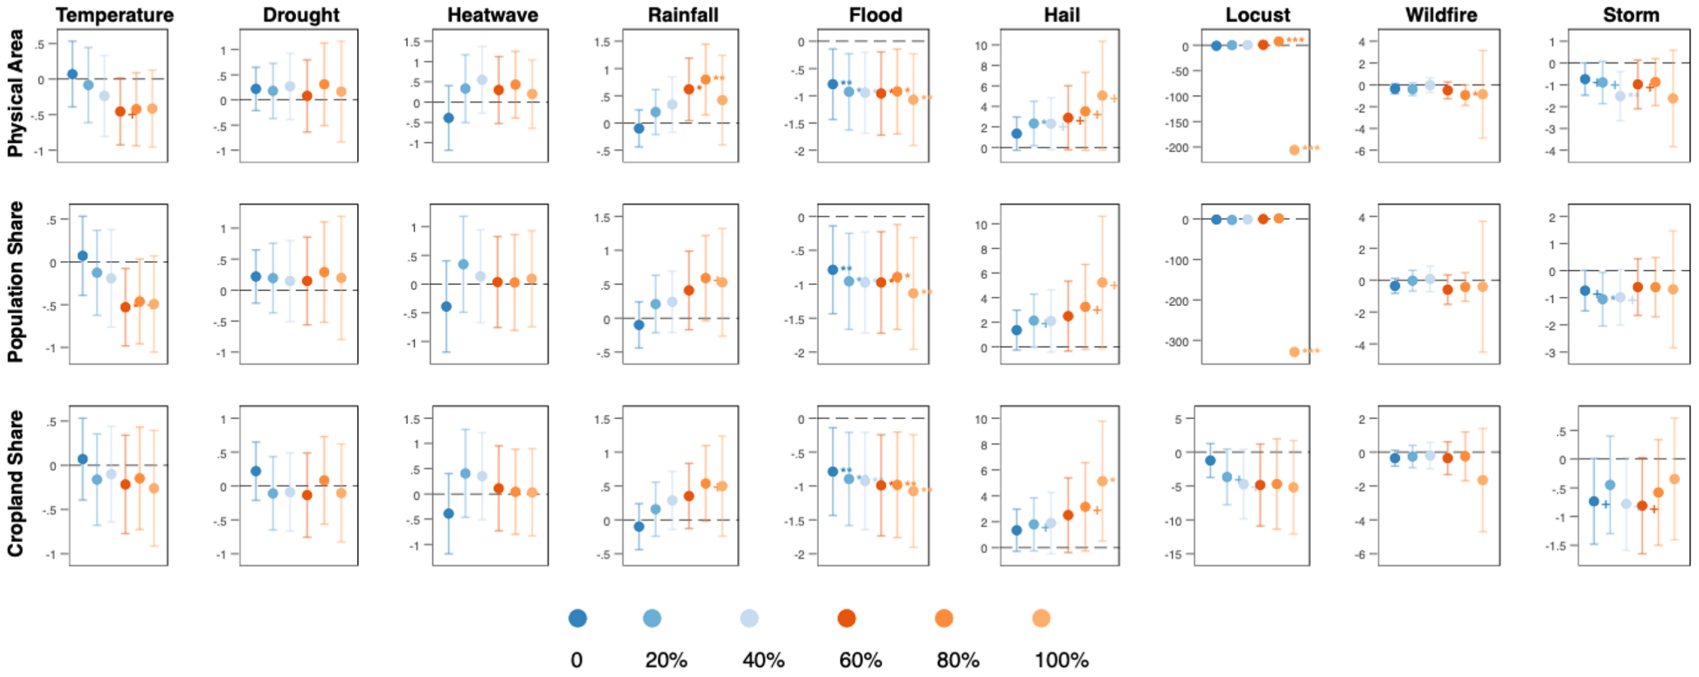

Fig. C.14. Robustness: Area Threshold

Note: *<0.1, **<0.05 and ***<0.01 p-value. The model estimates Equation 3.1. As mentioned in the methodology, treatments are restricted to affect at least 20% of the physical area of a district. Here, the threshold is changed from 0% to 100%. Also, alternative threshold criteria (population and cropland affected) are included.

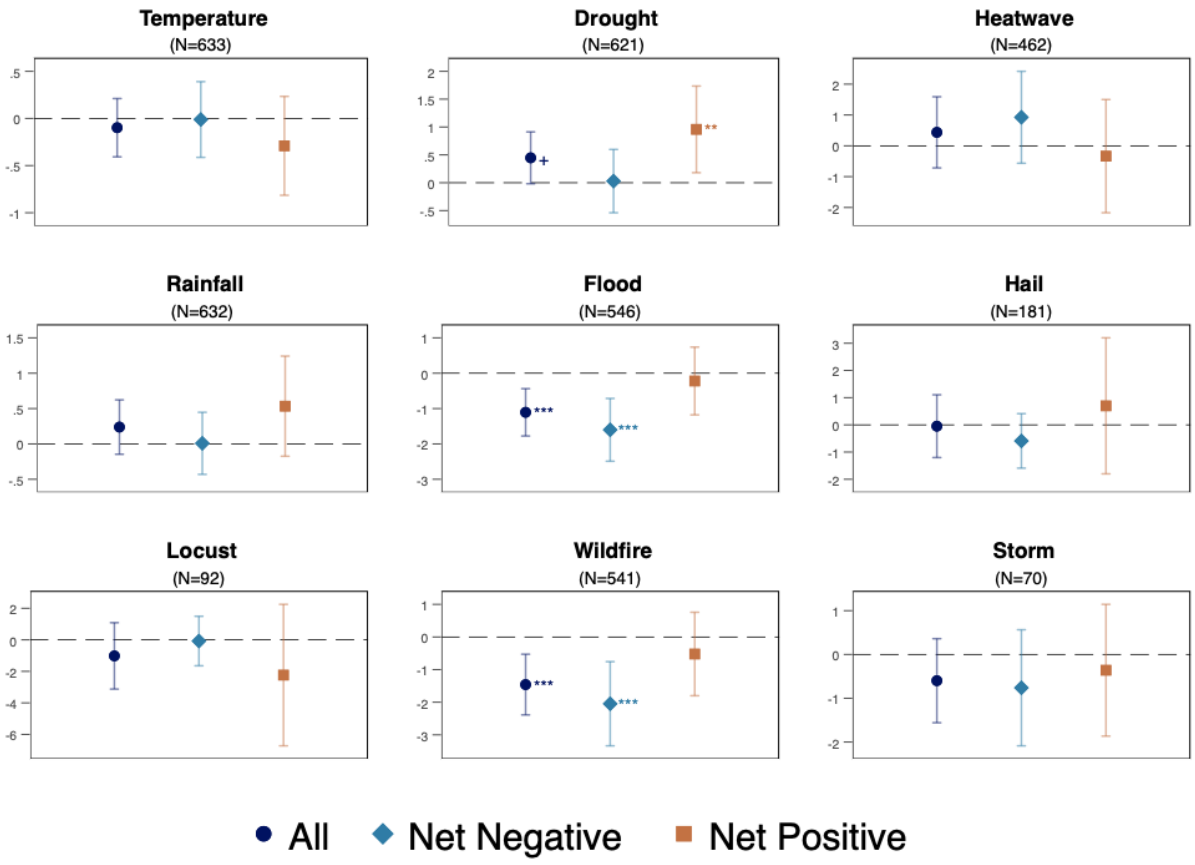

Fig. C.15. Robustness: Spatial Aggregation

Note: *<0.1, **<0.05 and ***<0.01 p-value. The model estimates Equation 3.1 but aggregates values at the province level.

C.5 Robustness of Extended Model

Concerns regarding Imperfect Multicollinearity There might be concerns of Equation 3.2 that variables of hazard interactions are added that are irrelevant or highly collinear to one another. Thus, we employ the Post-Double Selection Lasso (PDS-Lasso) method as proposed by Belloni, Chernozhukov, and Hansen (2014). It ensures that only hazards relevant to the outcome and treatment are included. In our case, treatment involves each separate coefficient for the 47 hazard combinations. Thus, the PDS-Lasso is re-run iteratively. We proceed in three steps: first, we select the variables most predictive of net population change using the Lasso method. Second, we select covariates that predict the hazard-combination in question. Third, and finally, a panel estimation is run using the treatment variable and the selected controls from the first two steps. In contrast to Equation 3.2 no region-year fixed effects are used as they are not compatible with the setup in the statistical software (STATA) used.

The results are shown Figure C.16. In general, the effects are broadly comparable, but cumulative events are more often significant (especially for droughts, heatwaves, floods and wildfires). The coefficients tend to be marginally larger than previously shown. The differences found might be due to a different set of included covariates or due to the omission of the region-year fixed effects. To make the results more comprehensible, we also create a difference plot between the coefficients of Equation 3.2 and the PDS-Lasso Model. The result is shown in Figure C.17. We find that if both are significant, the coefficients are fairly similar (close to the dotted line). In both cases, some coefficients are significant only in one model or the other. Four out of these are due to cumulative events that seem to become more significant in PDS-Lasso. Further investigation revealed that most of the significant ones were coefficients that, in the baseline, had p-values close to the 10% cut-off. Generally, there is a high degree of overlap between the direction and significance of the coefficients. This is especially true for the most impactful compound events. The PDS-Lasso detects some additional effects that might have been previously insignificant due to the inclusion of region-year fixed effects. This is especially due to the cumulative effects of temperature, drought, heatwave and flood. The differences in magnitude found might be due in part to the chosen variable selection, which might reduce bias in the final model. To conclude, the robustness of the direction of coefficients and the stability of the main compound coefficients, even after the selection of other relevant influencing variables, imply that the baseline effects are not systematically driven by the included variables.

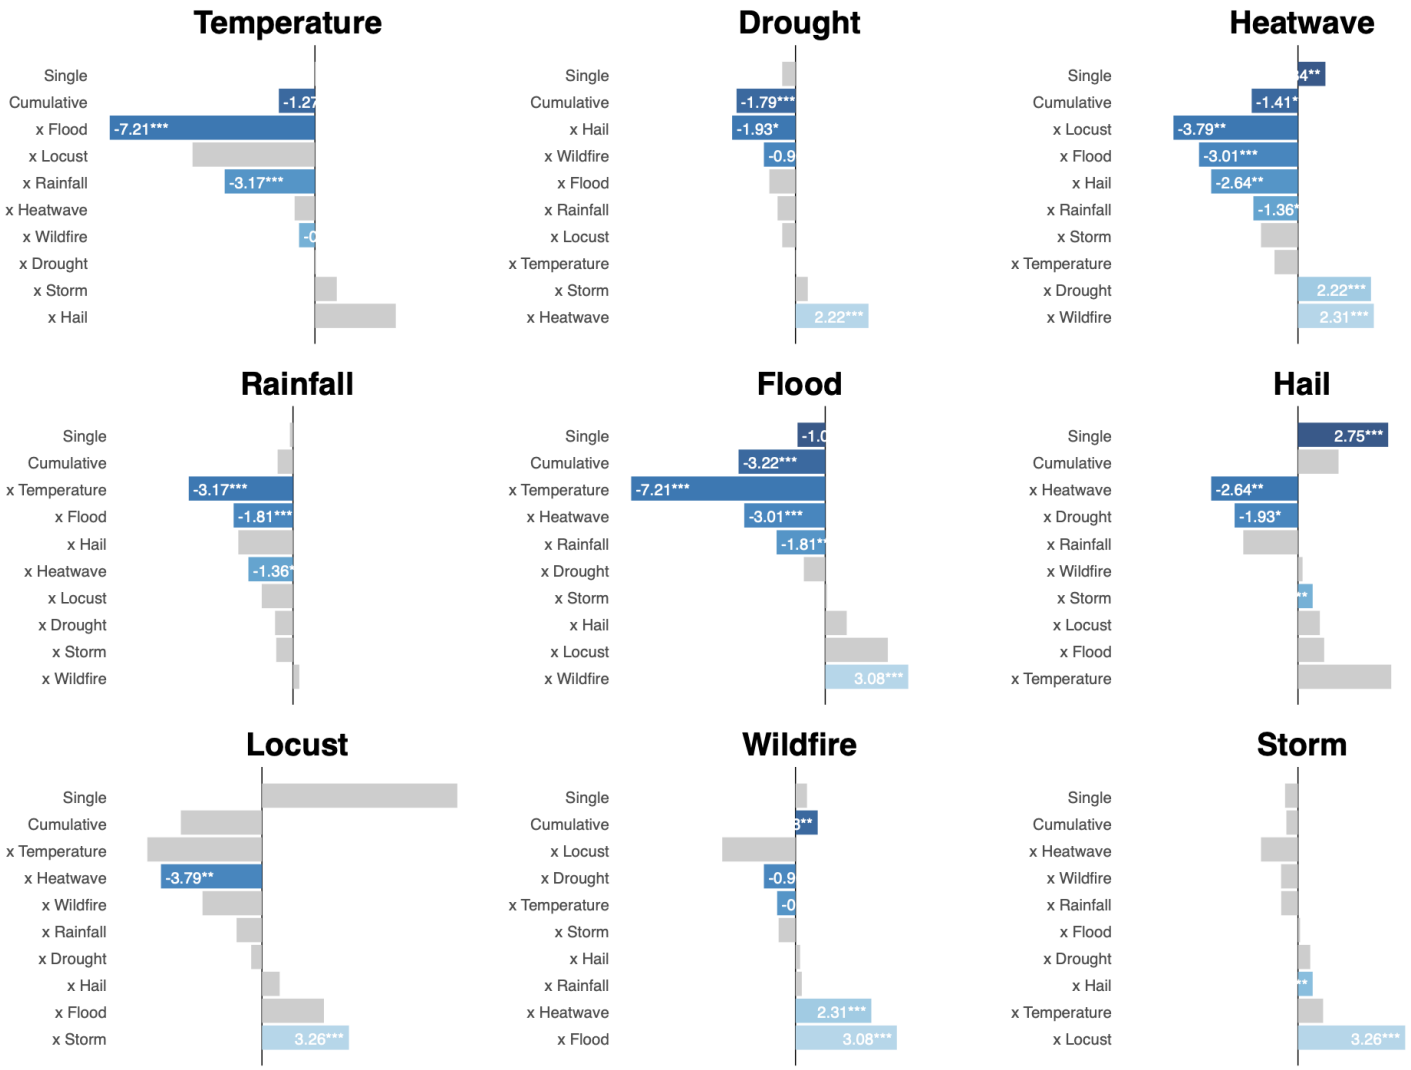

Fig. C.16. Robustness: Post-Double Selection Lasso

Note: *<0.1, **<0.05 and ***<0.01 p-value. The model yields similar results to Equation 3.2 but applies a post-double selection Lasso estimation, as described in Belloni, Chernozhukov, and Hansen (2014). Each coefficient employs a double selection of coefficients. Regressions include the same controls as in the baseline but no region-year fixed effects. Standard errors are clustered at the district level.

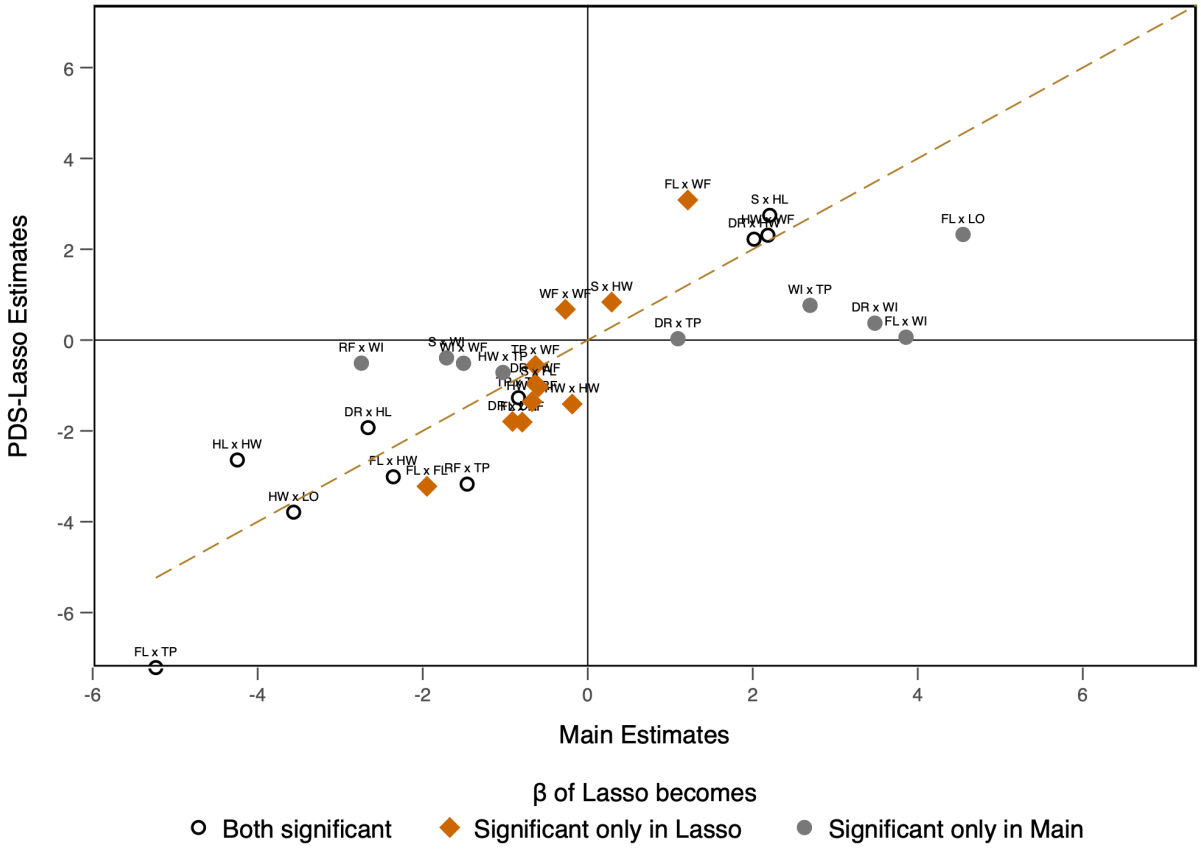

Fig. C.17. Robustness: Difference between PDS-Lasso and Baseline

Note: *<0.1, **<0.05 and ***<0.01 p-value.. The figure shows the difference between the estimates from PDS-Lasso and the baseline estimates. If both coefficients, PDS-Lasso and baseline, are insignificant, they are omitted from this figure. The closer both coefficients are to the 45-degree angle, the more similar the coefficients are. Abbreviations: T Temperature, HW Heatwave, DR Drought, P Precipitation, FL Flood, HL Hail, WF Wildfire and WI/S Storm.

Definition of Temporal Compounding The model shown before separates between single events, cumulative events, and compound events, where compound events are defined as two different types of hazards occurring in the same month. However, perhaps instead of overlapping within the same month, co-occurrence within the same year might also be important. Even if a flood and drought do not occur in the same month, their combined impact on net population change may be different from the impact of the individual hazard events. To test this, instead of using the sequence definition used in the main part of the paper, the same regression is used with the difference that "compound events" are now interactions within the year. So it is a less strict definition that allows a compound event to manifest at any point within the year. As long as both types of hazards occur in the same year, it is considered a compound event. Thus, the number of compound events substantially increases. The results are shown in Figure C.18. On average, many effects are similar (especially flood-temperature events), albeit smaller. Using interaction terms, many flood interactions show negative and significant effects. A new finding is that wildfire-flood interactions are strongly negative and significant; however, this is likely due to a relatively small sample of such events. For the effects that remain significant, changing the definition from within-month to within-year broadly keeps the sign across both models. Several events (flood-temperature, flood-storm) have a smaller magnitude of coefficients, suggesting that monthly compound events capture more immediate and severe impacts. The strongest effects remain significant, especially flood-temperature. However, a few additional effects are noted, notably flood-rainfall and flood-wildfire interactions that might be missed using monthly definitions. This highlights that people might respond not only to direct compound events, but also to cumulative or sequential shocks. To conclude, the chosen definition changes if a larger time window is used. However, the changed definition to within-year events does not invalidate the main results, given that the magnitude and direction of coefficients remain broadly stable. However, some events, such as floods and wildfires, might see a stronger response only if more time is between each individual event.

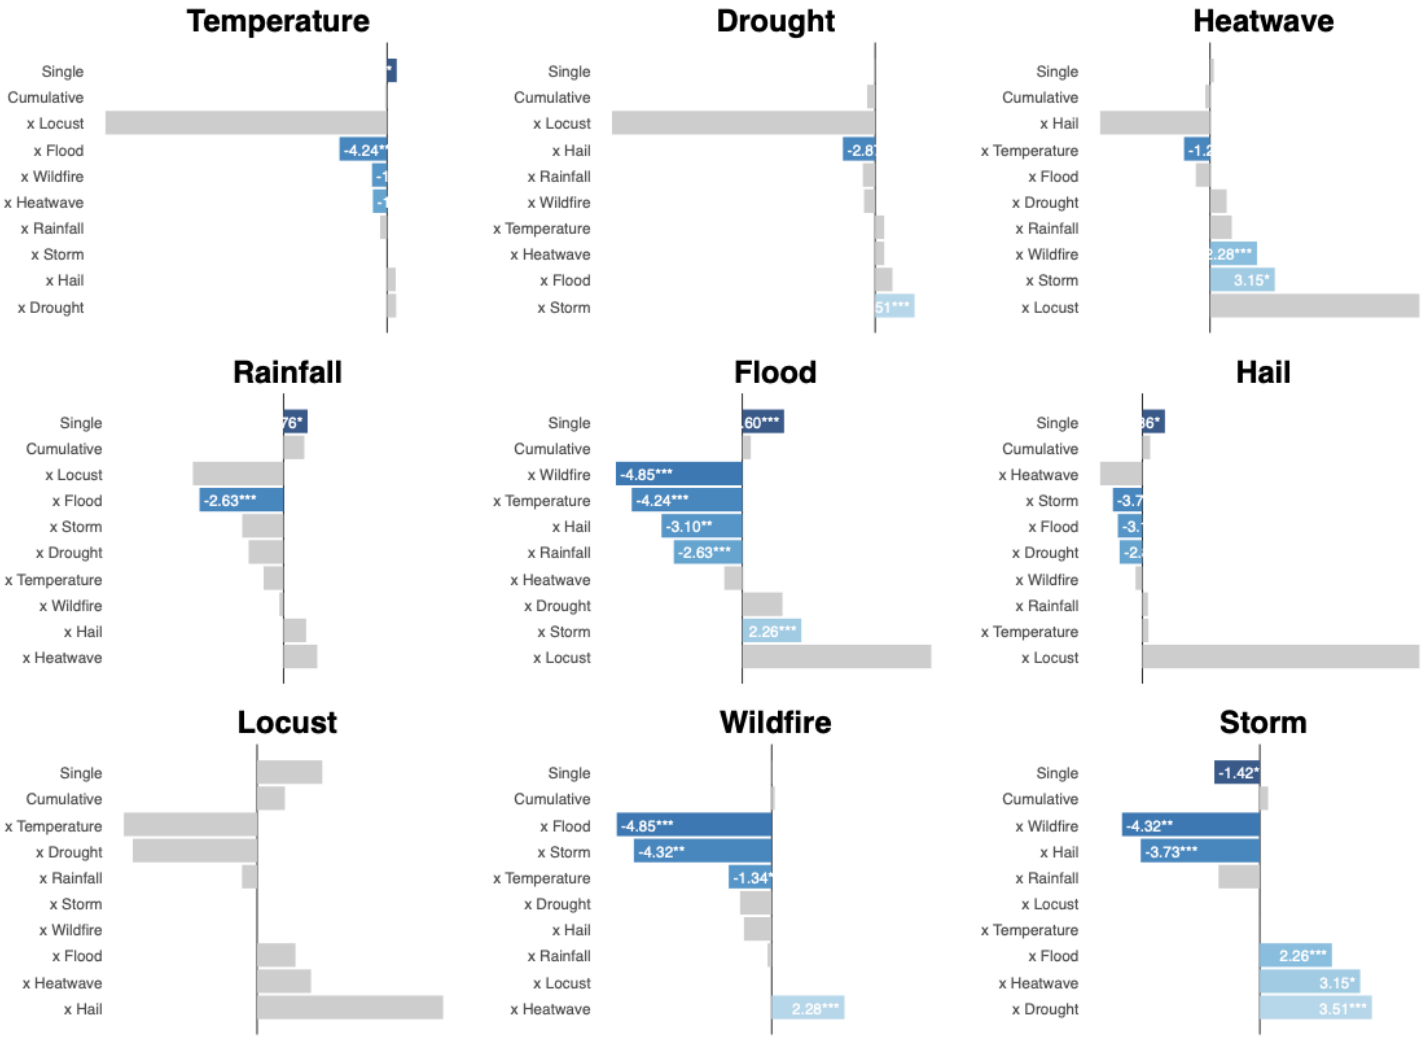

Fig. C.18. Robustness: Changing Compound Definition

Note: *<0.1, **<0.05 and ***<0.01 p-value. The figure shows the effects of Equation 3.2 with the difference that compound events are defined as interaction terms within the same year.

C.6 Illustration of the Vulnerability Index

In the heterogeneity section of the paper, the sample is divided into two groups: those with below-median vulnerability and those with above-median vulnerability. To our knowledge, no dataset exists at the sub-provincial level in Africa to specify climate vulnerability. To account for the vulnerability of a district, we rely on a downscaled version of the INFORM index, originally developed by the European Joint Research Center (Marin-Ferrer, Poljanšek, and Vernaccini, 2017). However, instead of replicating the entire indicator, we focus particularly on the "vulnerability component". Vulnerability is composed of two parts: the socio-economic dimension and the vulnerable group dimension. Socioeconomic factors themselves are composed of three parts: development deprivation, inequality, and aid dependency. We try to replicate this using the human development index from Sherman et al. (2023) for development deprivation. Inequality is approximated by a relative wealth index Chi et al. (2022). Aid dependency uses province-specific values for aid in USD received from AidData (2017). The vulnerable group is divided into uprooted people and other vulnerable groups. We do not have data for this, but we proxy the vulnerable population as the share of the population below the age of 15 and above 60 (as a proxy for age dependence) in each district, using data from WorldPop (2018). Each of these variables is then weighted according to the following formula:

$$\Omega_i = \frac{\mu_i - min(\mu)}{max(\mu) - min(\mu)} * 100$$

where μ is the observed value in district i, min(μ) and max(μ) are the minimum and maximum possible values of the variable and 100 is the scale to which we want to transform the variables. The index is then created as the sum of the individual weighted components, where similar to Marin-Ferrer, Poljanšek, and Vernaccini (2017), socioeconomic vulnerability with the weight of 0.5 (consists of development and deprivation, in form of HDI, has the weight 0.5, inequality, measured by the relative wealth index, has the weight 0.25, and aid dependency, here aid received as weight 0.25) and 0.5 of vulnerable groups, measured by population share below 15 or above 60. The final index ranges from 0 (little to no climate vulnerability) to 100 (highly vulnerable). A plotted version of the index is shown in Figure C.19.

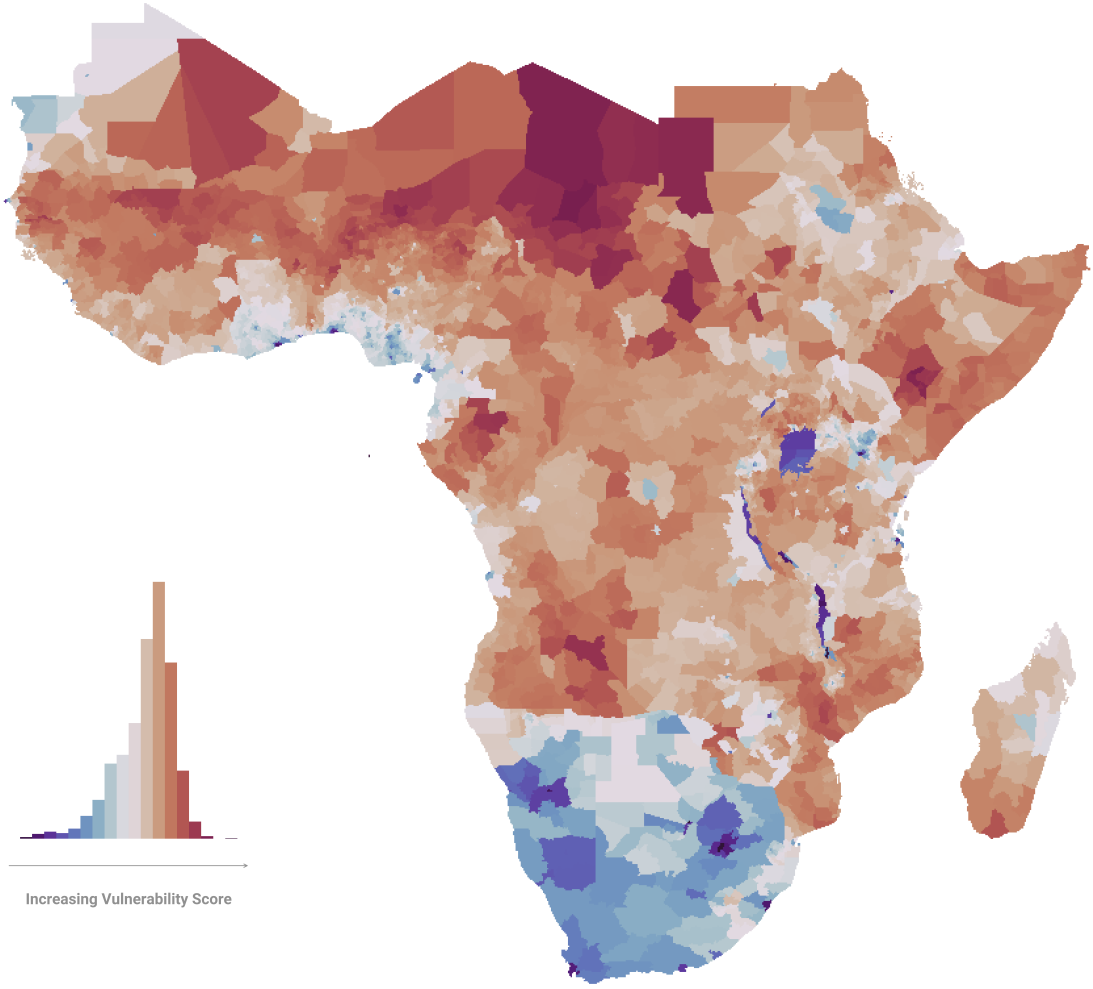

Fig. C.19. Visualization of Self-Composed Vulnerability Index

Note: The figure shows a self-composed vulnerability index based on the idea of the INFORM index. It combines the Human Development Index from Sherman et al. (2023), relative wealth from Chi et al. (2022), share of population below the age of 15 or above 60 (as a proxy for age dependence) from WorldPop (2018) and province-specific values for aid received from AidData (2017) into a unified index using min-max scaling. Grey areas are cells with no data available.

C.7 Additional Information on the Mechanisms

Creation of Outcomes

NDVI Anomalies NDVI anomalies computed are based on bimonthly NDVI estimates since 1982 from Li, Cao, Zhu, et al. (2023). Based on this, pixel-by-pixel anomalies are computed using the monthly means and standard deviations from 1982 to 2000, and then using a moving 18-year window. Then, an annual anomaly is computed as the average anomaly in a given year. However, this would capture the average anomaly by district when aggregated. To compute a district-specific measure, we use a cropland-weighted aggregate by year. To this end, we rely on annual cropland data from Winkler et al. (2021). The final measure measures the NDVI anomalies by district weighted by the fraction of cropland covered in each district. This is not an ideal measure given that NDVI cells might stretch over multiple land-use types not only cropland. Furthermore, originally, data could have been obtained on a daily 300m resolution, but the spatial and temporal aggregation causes a large loss of variation. The final annual NDVI anomaly at the district level on the croplands thus captures a smoothed generalized signal of cropland vegetation health by district rather than a high-resolution crop-specific stress dynamic. It may underestimate local and short-term deviations due to spatial and temporal aggregation.

Food Production The food production data on Cassava and Maize is taken from Lee, Anderson, et al. (2025) that cleaned and harmonized crop production data for 15 SSA countries. Data is provided at the district or provincial level, as well as by growing season and crop type. Our main data is annual data from the district level. We select Cassava and Maize for all areas where the growing and harvesting seasons are within the same year. Only production statistics from main growing seasons are kept (Main, Annual, Summer, Winter, Meher and Rice Season). That might underestimate true production during off-season planting, but it allows us to capture at least the major growing seasons. If data is only available at the provincial level, provincial statistics are distributed using the time-invariant harvested area by crop data from International Food Policy Research Institute (2020). Unfortunately, not every crop records the same number of seasons every year. To obtain a stable sample, we conduct multiple data cleaning steps, including manual validation of seasons covered and assessment of the bias. The bias is greatest for rice and wheat, where a significant amount of data is missing. Thus, in our final sample, only cassava and maize are included as they are also important staple crops in Africa. The final variable is logged to reduce the influence of production outliers.

Food Insecurity Unfortunately, food insecurity data is not available in all sub-Saharan African countries. We obtain data from 2009 until 2020 for Burkina Faso, Burundi, Cameroon, Chad, Congo, Democratic Republic of Congo, Ethiopia, Kenya, Uganda, Tanzania, Madagascar, Mali, Malawi, Mozambique, Niger, Nigeria, Somalia, South Sudan, Sudan and Uganda from Andrée et al. (2020). The data is based on the Famine Early Warning Systems Network on a tri-annual basis. The chosen measure is the Integrated Food Security Phase Classification (IPC), an index ranging from 0 to 5, where higher values indicate more severe famines. The data is already provided at the subnational level.

Gross Domestic Product Gross Domestic Product at the subnational level is difficult to obtain. Articles usually rely on nightlight as a proxy for economic activity. At the district level, nightlight falls short of taking into account the distribution of nightlight within the district if averages are used. We rely on a downscaled GDP at the provincial level, as presented in Zhang et al. (2024), which utilizes a deep learning model to estimate annual GDP at the provincial level. To distribute the provincial-level GDP to the district level, we use relative population shares within each province. Furthermore, we re-estimate the results using district-level estimates from 2012 onward, as provided by Rossi-Hansberg and Zhang (2025).

Changed Models

In the mechanism section, we rerun Equation 3.2 with different outcome variables. The variable transformations are described above. To measure the effect of different events on anomalies in NDVI, food insecurity, and crop production, the baseline model changes the temporal dimension from lagged treatment to treatment in the same period (switching t-1 to t). Apart from that the models are comparable. The reason for not using a lagged term is that we expect the effects of hazards to have an immediate impact on the results of NDVI and agricultural production. While it is also likely that effects carry-over to later periods, we use no lagged treatment here due to the following argument: if hazards have an immediate impact on agricultural health (NDVI anomalies), production and direct consumption (proxied by food insecurity) then effects on net population change a year later might be due to losses of agricultural income in the previous period. For GDP growth, the identical model as in Equation 3.2 (with t-1 in treatment) is used because we assume that hazards need time to manifest and influence GDP growth.

CHAPTER 4

Growth Under Rising Pressure: Weather Shocks in Sub-Saharan African Cities

This chapter is solely authored[1].

4.1 Introduction

African cities are expected to triple in population by 2050, compared to their size in 2018, and are estimated to house 22 percent of the world population by mid-century (UNDESA, 2019). These growing cities are major drivers of economic growth and have improved many socioeconomic outcomes, such as education and financial inclusion (OECD, United Nations Economic Commission for Africa, and African Development Bank, 2022). However, rapid expansion of cities, particularly those near the coasts, is likely to expose a growing number of residents and businesses to the impacts of climate change (Dodman et al., 2022; Rentschler, Avner, et al., 2023; Tuholske et al., 2021). Along with urbanization, one finds the sealing of the ground that might exacerbate flooding and heat exposure (Rentschler, Avner, et al., 2023; Founda and Santamouris, 2017). However, the effects of climate change are more complex than occasional exposure to a specific extreme event. By 2030, the likelihood of urban areas facing floods and droughts will have doubled compared to risk levels in 2000 (Güneralp, Güneralp, and Liu, 2015) and some cities will experience exposure to events beyond floods and droughts (Šakić Trogrlić et al., 2024). When multiple such events occur in a short time frame, often referred to as "compound events", their joint effects might be more severe than their individual effects (Zscheischler, Westra, et al., 2018). The increasing likelihood of co-occurrence of different extreme events and the rapid urban growth in Sub-Saharan Africa make it crucial to understand how these events shape urban growth.

[1] I would like to thank Guy Michaels, Clement Bosquet, Lisa Chauvet, Maxime Roche, Niklas Nordfors, Yanis Bekthi and the participants of the 14th European Meeting - Urban Economics Association, the IPSWD at Columbia university as well as internal research seminars at Paris 1 and LEO at University of Orleans for their valuable comments.

This article studies the effect of four types of climate-driven weather shocks on physical urban expansion in Sub-Saharan African cities between 2000 and 2019. Floods, heatwaves, storms and droughts are chosen as they are likely to change in the upcoming decades due to global warming (IPCC, 2021, Chapter 12). In contrast to previous studies, the four hazards are introduced simultaneously. This allows for investigating whether the effects found in the literature are driven by a specific hazard class or instead by the presence of another unaccounted-for hazard. Understanding how cities respond to extreme events is critical for several reasons. Although many studies have studied the relationship between climate change and urbanization (see, e.g. Henderson, Storeygard, and Deichmann, 2017; Helbling and Meierrieks, 2023), very few have focused on the effects beyond a specific hazard (mainly flood or drought) on urban growth. Given the increasing likelihood of co-occurring extreme events due to climate change, this understanding is particularly important for anticipating urban growth patterns. Extreme events may drive population movements (Hoffmann, Dimitrova, et al., 2020), intensifying pressure on growing cities but serving as an adaptive response to climate shocks. Yet, if these shocks exacerbate pre-existing vulnerabilities in cities, they might reduce welfare and induce additional humanitarian and economic pressure on cities. Lastly, many studies focus on the evolution of the population over time. The quality of this data is often limited since it relies on interpolated estimates. Using novel high-resolution building footprint data allows researchers to capture a clearer picture of extensions of urban areas.

To answer the question of how urban growth is changed after extreme events, we rely on a temporal panel model to compare the growth of 5 721 urban agglomerations in 44 Sub-Saharan countries between 2000 and 2019. After selecting urban agglomerations, as defined by OECD and SWAC (2020), the data is combined with high-resolution data on the growth of a measure called "impervious surfaces". This data, coming at a 30 m resolution, provides information on the annual evolution of human-modified areas within cities. Impervious surfaces capture, beyond building evolution, also evolution in infrastructure in general. It is especially important in climate impact studies, as it is correlated with a higher risk of heatwaves and urban flooding. Data on extreme events is drawn from Reinhardt (2024a) and includes information on floods, heatwaves, droughts and wind storms. The data is homogenized and converted into an annual count of shocks for each hazard that occurs in a city. However, the spatial scope of the events varies. Droughts have large-scale effects, while floods tend to be local. Thus, we repeat part of our analysis for a city-catchment area, which accounts

for the surrounding regions connected to the city. Then, while annual counts enable us to investigate whether an increasing number of shocks affects growth, we also examine the effects of a more "climatic" definition of compound events. Namely, we investigate whether our results change when accounting for different combinations of shocks that occur one after another within a year (e.g., a flood following a drought). Finally, in the article, we also provide insights on how much of the coefficients are driven by the history of exposure over the past eight years. Therefore, we rely on an event study estimator by Imai and Kim (2019).

Our main results paint a clear picture: experiencing floods, especially those induced by extreme precipitation, reduces urban expansion by around 3-9% in poorer and more rural cities. These effects are largely centered around Western Africa. In addition, these effects are amplified when floods follow heatwaves, which is a common type of "compound event". Rather than highlighting opportunities for sustainable growth, the findings are linked to undesirable side effects. We also document a contemporaneous increase in informal population, people returning to high-risk flood zones, as well as a slight increase in the likelihood of violent urban unrest. Heatwaves might have a small, negligible negative effect as they create potential disamenities in the city and reduce labor productivity (Lai et al., 2023). They might thus slow down urban expansion. Although some of our results suggest this, the effects are weak and negative. We find that most of the negative effects are caused by interactions with flood events. The effects of droughts have already been extensively studied, with numerous papers finding positive impacts on urban expansion (see, e.g. Chlouba, Mukim, and Zaveri, 2023). The mechanisms proposed in the literature focus on droughts as a factor that incentivizes rural-urban migration. Although little effect is observed when droughts are studied purely at the city level, accounting for the effects in catchment areas of cities, we document an increase in urban expansion of around 3% (or 0.12 percentage points). In partial contrast to the literature (e.g. Van der Borght et al., 2025), however, we find that these effects originate mainly in medium-sized rural cities. These effects become negative once the history of droughts is taken into account, suggesting that treatment history matters. Finally, storms, although they show strongly positive effects, should be regarded with caution, since only small fractions of African cities are exposed to this hazard. However, relatively wealthy cities experience an increase in urban expansion after damaging windstorms, highlighting that they can induce population movements. Thus, combining the results, we find that shocks, especially droughts, tend to reveal urbanization as an adaptive response to surrounding areas. However, the shocks stud-

ied also increase vulnerabilities within cities. From a policy perspective, floods have a major influence on urban expansion in Sub-Saharan Africa. They add pressure to already poorer cities and exacerbate living conditions. Finally, from an empirical perspective, we emphasize that in addition to the history of hazard exposure and the spatial dimension of the shock, researchers should consider sequences of different hazard events. This approach might help us better understand whether the hazard alone or a combination of hazards drives the results.

This article contributes to several streams of literature. First, the article contributes to the literature on the effects of shocks on urban growth and recovery from shocks. This topic has been extensively studied in urban economics in the context of war (Davis and Weinstein, 2002; Bosker et al., 2007), human-caused events such as the Boston Fire of 1872 (Hornbeck and Keniston, 2017) or in the aftermath of particularly costly extreme events such as Hurricane Katrina in 2005 (Vigdor, 2008). The paper contributes to this discussion by examining how extreme events influence urban growth and expansion. Instead of focusing solely on a single extreme event, such as a fire or a hurricane, we extend this literature by studying, in a high-resolution setting, how cumulative rather than single shocks might influence growth behavior.
Second, it contributes to the literature on the urban growth-climate change nexus. In recent years, an increasing number of studies have shifted their attention from individual cities or events to the effects of climate change on urban growth. An entire chapter of the IPCC is dedicated to this topic (Dodman et al., 2022). Many notable examples in economics study the effects of specific hazards or joint hazards such as temperature and precipitation (Castells-Quintana, Krause, and McDermott, 2021; Helbling and Meierrieks, 2023; Colmer, 2021), droughts (Henderson, Storeygard, and Deichmann, 2017; Chlouba, Mukim, and Zaveri, 2023; Van der Borght et al., 2025) or floods (Gandhi et al., 2022; Kocornik-Mina et al., 2020; Magontier and Martinez, 2023). Many of these studies are at the city level. The closest related to this study are the works of Magontier and Martinez (2023) and Kocornik-Mina et al. (2020) for floods and Van der Borght et al. (2025) for droughts. Magontier and Martinez (2023) focus on the effects that extreme floods, caused by rainfall shocks, have on urban development in Spain. Applying an event-study approach, they document that floods have led to the vertical expansion of cities. This expansion has generally contributed positively to structural change and increased amenities in affected cities. Kocornik-Mina et al. (2020), in contrast, focus on a more global sample of cities and document heterogeneity depending on different elevation levels within cities. Affected areas do not appear to adapt by shifting

economic activity away from risky areas, except in areas that are relatively recent. In contrast to Magontier and Martinez (2023), our study focuses on a broad scale of developing countries, where the effects might be fundamentally different. Contrary to Kocornik-Mina et al. (2020), we actively avoid relying on nightlight data, as older data, such as DMSP, tends to be biased in post-disaster studies (Gibson, Jiang, et al., 2024). Instead, we use physical urban expansion. This metric, changing more slowly, is more robust against the large fluctuations of nightlight. The article by Van der Borght et al. (2025) is the closest related to this study, as it examines the effects of urban expansion between 1984 and 2015 following severe drought events. They document significant urban expansion in metropoles of affected countries, where, especially in high-flood-risk areas, build-up appears to increase. While relying on the same input to define urban areas as this study, we extend their work in three ways: we account for other hazards than only droughts, we study the temporal dimension accounting for the history of droughts in a city and we do not exclude small cities as they might serve as an important piece in the impact dimension (if displaced persons tend to settle first in nearby towns due to lower displacement costs). In contrast to existing literature, we document the importance of accounting for the co-occurrence of other hazards, either by including them simultaneously in the model or by testing for specific interactions. Furthermore, our article discusses the spatial scale of hazard events and emphasizes the importance of going beyond continental heterogeneity analysis, which might overlook within-continent hazard variation due to variation in atmospheric conditions.
Finally, we attempt to take a more holistic approach by including multiple types of hazards in the analysis. Although the importance of analyzing connected and compound events is well-documented (Zscheischler, Westra, et al., 2018), an emerging body of research examines their impact on particular cities (Šakić Trogrlić et al., 2024) or specific zones, including coastal and valley regions (Dodman et al., 2022). Some also focus on specific hazard combinations in a specific country (Ghanbari et al., 2023). Despite the simplified model used here and its failure to account for the full complexity of compound events, the paper makes a contribution by studying, in its extension, the effects of combining sequences of hazards. In summary, previous literature has primarily focused on general shocks or how specific weather shocks drive urban expansion, and has documented heterogeneous growth responses. Little attention is given to moving beyond single hazards and accounting for the co-occurrence of other hazards. This article aims to fill the gap by providing a broader understanding of how different weather shocks impact urban expansion and whether their temporal and spatial dynamics, as well as their interactions, matter.

The paper is structured as follows: first, the data (section 4.2) and the empirical framework (section 4.3) are presented. Next, the results, their heterogeneity and robustness are discussed in section 4.4.1. Section 4.4.2 adds extensions to the core model by accounting for compound events, different spatial aggregation and finally accounting for treatment history. The results are placed in literature and the question is asked if the effects found improve or deteriorate the condition of the city in section 4.4.3. Finally, the paper is concluded.

4.2 Data

This article relies on a large variety of data sources. In the following, the main variables are described and the way in which the panel set is constructed is explained. A general description of how data is generated is shown in Figure 4.1.

Urban Areas At first, a definition of an urban area is needed. In this study, we rely on urban agglomerations as defined by the Africapolis by OECD and SWAC (2020), similar to the article by Van der Borght et al. (2025). This dataset applies spatial criteria (area where buildings are less than 200m apart) and population criteria (at least 10,000 people) to define an "urban agglomeration". For the Sub-Saharan sample around 5 721 agglomerations are included. In contrast to Van der Borght et al. (2025), we do not apply exclusion criteria to the sample. The study aims to include smaller, rapidly growing towns that can serve as initial destinations for displaced populations from nearby areas. We aim to thereby obtain a more complete picture of urban development in Sub-Saharan Africa. In the robustness part, we test how changing the sample of cities changes the results. Figure 4.1 shows, in the first column of panel A, the agglomerations included in Nigeria, and in panel B, the shape of a representative agglomeration. Throughout the article, we refer to these urban agglomerations as "cities", while acknowledging that the differences in size between them may be relevant. A possible alternative could have been the use of the Global Human Settlement Layer (GHSL) (Pesaresi, 2023). However, we prefer to rely on Africapolis, given that it relies on a standardized definition of urban areas for the African context, whereas GHSL is designed to be globally applicable, thus losing some scope for African cities.

4.2.1 The Dependent Variable

Physical Urban Expansion To measure the physical expansion of urban areas, we rely on the evolution of impervious surfaces at a resolution of 30 m from the Global Imper-

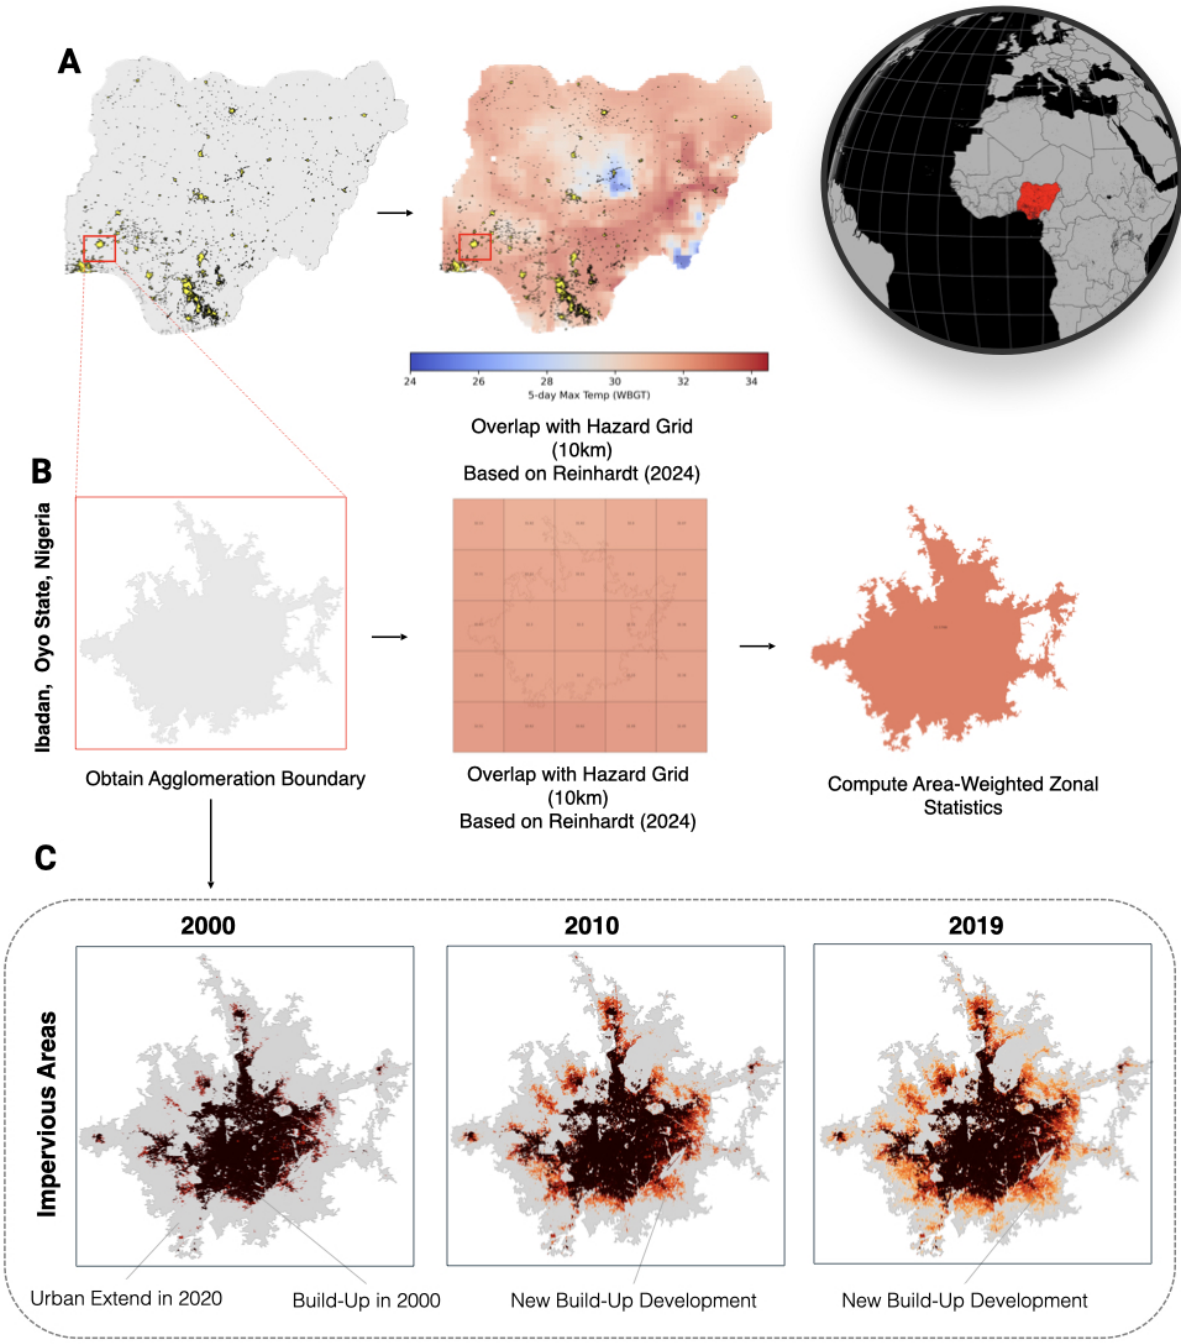

Fig. 4.1. Illustration of Data Generation: Nigeria

Note: Panel A displays Nigeria with its urban agglomerations, as defined by OECD and SWAC (2020), alongside the added hazard grid from Reinhardt (2024a). Panel B focuses on a single agglomeration—Ibadan in Oyo State—showing its overlap with the 10x10 km grid and the computation of weighted zonal statistics (in this case, for the 5-day maximum Wet Bulb Globe Temperature in March 2010). Finally, Panel C presents the raw data alongside the assigned data for the outcome variable—physical urban growth. Red and yellow areas are areas added after 2000. The grey background is the urban extent of the city based on Africapolis.

vious Surface Area (GISA) dataset (Ren et al., 2025). This data comes as a raster where each pixel is indicated as built or not built, as well as the year when it was built. Impervious surfaces (e.g streets, buildings, and parking lots) are human-made structures that prevent water from being absorbed into the ground. They are indicative of physical development. More precisely, it measures the horizontal physical expansion of a city, allowing cities to grow in space rather than height. If poorly managed, they are also strong predictors of future vulnerability to hazards such as flooding (Feng, Zhang, and Bourke, 2021) and extreme heat (Huang, Masselot, et al., 2023). Beyond their environmental significance, the expansion of impervious surfaces can also be interpreted through an economic lens. Urban growth may result from productive economic activity that leads to urban expansion or from booming, unregulated informal settlements, often triggered by displacement and a lack of affordable housing. Thus, in this study, urbanization itself is considered neutral. If it leads to economic growth, it would be considered positive. However, if it is accompanied by increased informality and more people living in vulnerable areas, urbanization would be negative. Although not the key focus of this study, we assess whether the effects found are rather positive or negative in the "Discussion" section of this paper.
In the literature, one usually finds either nightlight or the WSFevo dataset[2] (German Aerospace Center (DLR), 2024; Marconcini et al., 2020) to study urban expansion. We include WSFevo as a robustness check due to its widespread use. However, the selected GISA dataset offers two key advantages: longer temporal coverage (2000-2019) and more climate-relevant growth (since it includes infrastructure). Nightlight data is not reliable for hazard studies. Changes in luminosity reflect not only physical urban growth, but also economic activity and access to electricity. Furthermore, the type of satellite sensor used in nightlight studies might distort the impacts (Gibson, Jiang, et al., 2024).

Using 2020 urban agglomeration boundaries from Africapolis, we create bounding boxes around each city to measure annual impervious surface area from 2000 to 2019. As shown in panel C of Figure 4.1 for Ibadan, Nigeria, we use the 2020 functional boundary as the maximum extent of the city. We then calculate the number of built-up pixels each year, with brighter orange indicating later expansion. Each 30×30m pixel equals 0.0009 km², allowing conversion to a total human-modified area in each city.

[2] See, for instance, Van der Borght et al. (2025), Chlouba, Mukim, and Zaveri (2023), and Rentschler, Avner, et al. (2023) for applications.

4.2.2 The Independent Variables

The hazard data is predominantly drawn from Reinhardt (2024a), a dataset that combines pre-existing datasets from 2000 to 2019 into a unified global grid. The data is reported on a 10x10 km grid for each month, primarily in terms of hazard intensities. For this study, data on floods, droughts, heatwaves and extreme wind/tropical cyclones are extracted. An illustration for the heatwave metric, Wet Bulb Global Temperature, for the case of Nigeria can be seen in Panel A of Figure 4.1. As the outcome variable is annual and most cities are substantially smaller than the grid cells (averaging about one-fourth their size), some transformation is needed. This transformation is necessary to align the data with an annual panel at the city level. All described hazards are characterized by monthly continuous intensity functions[3]. First, a monthly area-weighted average or maximum of the hazard intensity is calculated at the city level to account for potential overlap across multiple 10×10 km grid cells. An illustration of this process can be seen in panel B of Figure 4.1. This value is then converted into a binary shock indicator (shock / no shock) using hazard-specific thresholds. The annual city-level shock is defined as the number of months the city transitions from non-treated to treated status[4]. A mathematical representation of the procedure is shown in Appendix D.3. The following paragraphs briefly describe the intensity measure for each of the four weather shocks[5]. They also explain the threshold used to define a shock for each hazard. The statistics and original data sources are presented in Table 4.7.

Floods Floods are well studied in the context of urbanization (see, e.g. Kocornik-Mina et al., 2020; Rentschler, Avner, et al., 2023; Magontier and Martinez, 2023). If older, less productive assets, such as machinery or buildings, are destroyed in catchment areas and replaced with newer, more productive ones, floods could contribute to economic growth in parts. In literature, this concept is sometimes referred to as "constructive destruction" (Hallegatte and Dumas, 2009; Hsiang and Jina, 2014). However, floods within cities can cause economic disruption (Kocornik-Mina et al., 2020), evasion of flooded areas, especially near the coast (Bukvic and Barnett, 2023), or zoning restrictions that reduce urban growth. We expect a reduction in urban expansion, as floods

[3] With the exception of flood, which has a categorical value range

[4] That means if a flood occurred from February-March and July-August, two events will be recorded. Certainly, the duration of an event might also matter, especially for floods and droughts. Accounting for the duration is assessed in the robustness section of this paper

[5] We refrain in this article from calling the observed shocks "climate" shocks, albeit they are heavily influenced by climatic change. Our aim is to emphasize their short-term, high-frequency impacts consistent with the framework of climate-driven weather shocks (Kolstad and Moore, 2020)

are notoriously destructive and could reduce incentives to grow. Although often used in the literature, rainfall itself is not a sufficient predictor of floods, as floods can originate from upstream river basins. We use the Dartmouth Flood Observatory (DFO) Index, which reflects both flood magnitude and likelihood. Since floods are common across much of Africa, we define "extreme" events as those with a return period of approximately 1 in 50 years (DFO Index > 1.5). For robustness, all recorded flood events are also included. Furthermore, in the robustness section, we distinguish by type of flood, assuming that a monthly 2-standard deviation from historical means in precipitation might be indicative of a *pluvial* flood. Similarly, a two-standard deviation from historical means over a 3-day period of maximum river discharge could be indicative of a *riverine* flood[6]. Finally, since not all parts of an urban area are equally impacted by a flood, we add information on flood risk from from Rentschler (2024), originally derived from Fathom Global, and coastal flood risk from Vousdoukas et al. (2018) and Vousdoukas (2023). Figure 4.3 visualizes on a province level the average total exposure of cities to floods in Africa for the period 2000 to 2019. Most floods are centered around Eastern Africa, Madagascar and parts of Nigeria.

Heatwave Heatwaves, though less clear than floods, could influence urban growth by reducing manual labor productivity (Zhao et al., 2021). In addition, they create heat-related discomfort and risk in densely constructed urban areas (Huang, Masselot, et al., 2023), reducing incentives to move into cities. Thus, we expect heatwaves to have a minor or no clear effect on urban growth. For heatwaves, the Wet-Bulb Globe Temperature (WBGT) is chosen as a metric, in particular the average WBGT over a 5-day window in a given month. Then the 5-day monthly maxima are taken. Using a fixed threshold does not account sufficiently for geographic variation. Some parts of Western Africa experience very high WBGT values almost every year. Thus, in addition, we compute a standardized score based on historical daily WBGT temperatures from 1970 to 1999 to calculate heatwave anomalies. A heatwave month is a month in which a 2 standard deviation anomaly occurred. Furthermore, wet-bulb temperatures must have exceeded at least 33 °C, a threshold that is considered extreme for the human body. Figure 4.3 shows that despite accounting for these extreme values, Western Africa is still the region where most heatwave events are recorded.

Drought Prolonged drought, as a slow-onset event, can be damaging for agricultural

[6] The data is drawn from the Global Flood Awareness System (Joint Research Center and Copernicus Emergency Management Service, 2019) on daily river discharge in m^3 per second. They are standardized on monthly means for the covered period (2000-2020).

production (e.g. Santini et al., 2022) and therefore lead to food insecurity. They could accelerate urbanization (see e.g.Van der Borght et al., 2025; Chlouba, Mukim, and Zaveri, 2023) through rural-urban migration, especially if droughts occur in catchment areas of urban areas. We expect that drought may lead to increased urban expansion through the indirect, non-observed channel of rural-urban migration. For drought, the paper relies on the Standardized Precipitation Evapotranspiration Index (SPEI) for a 3-month (SPEI-3) and 12-month (SPEI-12) period. Having the monthly intensities on a city level, an "extreme drought" is defined as a situation where both SPEI-03 and SPEI-12 are below a value of -1.5 in the same month, commonly referred to as severely dry environments. As droughts in some regions are fairly common, both conditions must be met. A similar approach has been used in the past to define "compound droughts" (see e.g. Wu et al., 2022). Given that droughts are fairly large-scale events, city-level treatment is likely to also capture similar drought conditions in the catchment area of a city. Figure 4.3 shows that most parts of Subs-Saharan Africa have seen at least a single drought over the period covered[7]. In our sample, this hazard appears to be one of the most common, as becomes clear in the descriptive statistics discussed below.

Storm Damaging winds, as a rapid-onset type of event, tend to damage the infrastructure and cause economic losses (Hsiang and Jina, 2014; Heinen, Khadan, and Strobl, 2019; Elliott et al., 2019). They might lead to unplanned expansion of cities in response to rebuilding destroyed housing. They could also lead to increased urban growth in cities away from the coast, as many storms occur near the coast. We therefore expect that windstorms have a positive effect on urban expansion. In this study, extreme wind and wind damage can be caused by wind speed anomalies or cyclones. Thus, the wind intensity combines the highest wind speed in a month recorded by tropical cyclones or 10m wind gusts. To be considered "extreme", a wind speed of at least 25 meters per second is required. This threshold is used, for example, by the European Severe Storm Laboratory (ESSL) (European Severe Storms Laboratory, 2014). However, Figure 4.3 shows that storms exceeding this threshold are relatively rare in Sub-Saharan Africa. They are mainly concentrated in Madagascar and a few provinces in southern Africa. Although the sample itself might be relatively small, we have included storm as a fourth category nonetheless. This decision was made in light of the expected increase in the intensity of future storms.

[7] One might wonder why countries such as Mali and Niger see no drought events. In such hyper-arid and consistently dry climates, it is more challenging to trigger a "compound" drought, as defined here. This is more common in regions with more rainfall variability

Other Controls

As main controls, similar to Van der Borght et al. (2025), we use country-level Gross Domestic Product (GDP) per capita and share of employment in agriculture over total employment from the World Bank (through Azevedo, 2011[8]). GDP per capita and employment in agriculture aim to capture important time-varying factors that could affect the response to extreme events. Poorer or more agriculturally dependent countries might be more vulnerable to these weather shocks (Meierrieks, 2021; Cattaneo, Beine, et al., 2019). Thus, we control for time-varying country characteristics. Furthermore, we include a population estimates from the global gridded population dataset (GlobPOP) by Liu, Cao, et al. (2024a). Urban expansion that excludes population estimates overlooks a significant aspect of what drives urbanization. We refrain from defining urbanization by relying exclusively on population estimates, given their interpolated nature. However, the GlobPop data is documented to outperform other existing population estimated datasets and is therefore used in some parts of that article. We expect that population trends follow broadly urban expansion. Finally, wealth estimates are taken from Lee and Braithwaite (2022) and estimate a time-invariant index of the absolute wealth of a location. This follows a similar reasoning as accounting for GDP per capita on a country-level and aims to allow us to distinguish between poorer and richer cities in relative terms within each country. All other data is described in the respective sections.

4.2.3 Descriptive Statistics and Limitation of the Data

The descriptive statistics of this article are shown in Table 4.8. Originally, the sample for Sub-Saharan Africa included 5 821 cities, of which 5 721 are included. The excluded cities show growth rates in urban expansion of above 150%. They were manually validated as outliers and therefore excluded[9]. The final sample covers 5 721 cities for the period 2000 until 2019. Out of this sample, 59 can be considered metropolitan areas[10] as population and economic centers of a country. The average city has only about seven square kilometers in impervious area. Two things can be taken from this: many of the

[8] For Eritrea, missing GDP per capita data after 2012 was supplemented using data from the IMF World Economic Outlook (accessed May 5, 2025).

[9] Growth rates of 100% are rare in our sample, but they can occur when relatively small urban agglomerations expand rapidly. Because these areas start from a low base, even modest absolute increases in urban extent can result in very high growth rates.

[10] This definition is provided by OECD and SWAC (2020).

our sample are small towns. Furthermore, with a standard deviation of 42 square km, there is a large heterogeneity in urban size in our sample. Columns (2) to (5) divide the sample into African subregions[11]. Most small cities are found in eastern and central Africa. Across the entire sample and all regions, cities exhibit annual growth rates of approximately 3%. These rates are computed as first log-differences. The by far most common hazards are droughts. An average city sees around 0.1 drought events a year or roughly one drought every ten years. Similarly, floods are also common. Heatwaves, as discussed earlier, are most commonly found in western Africa. In terms of control variables, most cities lie in countries where non-negligible shares of the population are employed in agriculture. In general, the descriptive statistics suggest significant variation between subregions. Thus, we test heterogeneity by subregion. Furthermore, when we count the number of cities by country, Nigeria represents one-fifth of the cities in our sample (see the lower panel in Figure D.2). This proportion mirrors the demographics of the real world, where approximately one in five people in Sub-Saharan Africa lives in Nigeria. However, this motivates robustness checks of sample dependence, which will be included in the robustness section of this article. A large part of cities are at high flood risk. The included towns of central Africa show, on average, the lowest wealth but the highest share of agricultural employment at the country level. In terms of shock computation, most cities overlap between two and three 10x10km cells, so city-level exposure values are based on averages of weighted statistics of two 10x10km cells by Reinhardt (2024a).

There are some notable data limitations. The measure of urban expansion, impervious surfaces, measures horizontal urban expansion rather than vertical expansion. So, if the city cores grow in height, this growth will not be captured. Given the focus on many smaller urban areas, horizontal expansion is likely the main driver of urban growth, in contrast to studies that focus on the effects in already densely populated regions (Magontier and Martinez, 2023). The hazard data is also limited. Converting hazard intensities into binary dummies may oversimplify the complexity of hazard events. There are two main concerns with this method: first, having a single dummy assumes homogeneous treatment effects and excludes the possibility of nonlinearity. Given that the thresholds are set to rather extreme values, the binary indicator may fail to capture the variation in the hazard intensity just below the threshold. This could potentially lead to underestimating weaker yet still impactful events. This study focuses on fairly high hazard values, making this threshold choice both interpretable ("one

[11] The division is illustrated in Appendix Figure D.2.

additional event") and consistent with parts of the existing literature[12]. For example, flood events are frequently operationalized as binary indicators, as seen in Kocornik-Mina et al. (2020). A second concern relates to how "extreme" is defined. According to the Intergovernmental Panel on Climate Change (IPCC) (Seneviratne, Nicholls, et al., 2012), all four hazard categories considered in this study are influenced by global warming. Consequently, these hazards may become more or less extreme. However, since the thresholds for defining extreme events are set retrospectively, they are based on a relatively short 20-year period. It is unlikely that substantial shifts in the underlying distributions of each hazard would be detectable within this time frame.

4.3 Empirical Framework

This study attempts to assess the effects of the four major extreme weather hazards on horizontal physical urban extension in Sub-Saharan Africa. We focus particularly on changes in urban expansion growth rates in the years following extreme events. In contrast to previous work, we account for the presence of other climatic shocks by explicitly including them in the regression model.

4.3.1 Specification

The main specification of this article is a panel regression model of the following form:

$$\Delta \log(Y_{icrt}) = \alpha + \sum_{k=1}^{4} \beta_k \cdot \text{Haz}^{(k)}_{icrt-1} + \gamma X_{ct-1} + \lambda_t + \lambda_i + \lambda_{r \times t} + \epsilon_{it} \tag{4.1}$$

The main outcome variable in this model is $\Delta \log(Y_{icrt})$, which measures the logged first-difference of the area in m^2 of a city that is covered with impervious surfaces, so human-modified land. Cities (i) are located in country c that again is located in sub-region r (Western, Eastern, Central or Southern Africa) and recorded in year t. The $\beta_k \cdot \text{Haz}^{(k)}_{icrt-1}$ represents the dot product between the regression coefficients β_k for each hazard type k (flood, heatwave, drought, storm) and the corresponding components of the hazard vector Haz_{icrt-1}, which includes all hazard variables lagged by one year. The four hazards, as described in the previous section, measure the number of events

[12] This concerns particularly articles relying on event studies that often reformulate treatment intensities into binary dummies to fit into treatment models

of this type in a year[13]. Hazards are therefore considered cumulative and additive, based on the total number of events per year. They are lagged under the assumption that shocks require time to influence urban growth. As controls X, similar to Van der Borght et al. (2025), the share of agricultural employment over total employment and GDP per capita in a country are included as lagged terms. They account for important factors that vary in time and could influence the response to urban growth. λ_t control for common global shocks such as commodity price shocks, pandemics, but also common trends in weather (see Kolstad and Moore, 2020). λ_i control for city-specific time-invariant effects such as distance from the coastline, average elevation, accessibility by road, and remoteness. Finally, we include region-year fixed effects ($\lambda_{r\times t}$) to capture region-specific trends, particularly those driven by large-scale climate teleconnections (e.g., El Niño Southern Oscillation, Madden Julian Oscillation), which tend to affect entire sub-regions rather than individual localities. These trends might heavily influence the occurrence of hazards and serve as omitted variables. The countries by subregions are shown in Table D.1. The standard error is clustered at the city-level[14].

4.3.2 Identification Strategy

The baseline coefficients (β_k) measure the short-term response in the year following an event to changes in physical urban growth resulting from an additional hazard event k. Identification comes from the temporal random variation in the hazard occurrence between and within cities. Following Kolstad and Moore (2020), the use of this variation in panel models describes a short-run response. In this short-run scenario, we expect little to no adjustments in beliefs or adjustments in investment, in contrast to longer-run effects of climate variation ("shift in distribution"). Given the calibration of the shock to extreme values, however, we expect that shocks nonetheless will create non-negligible effects. Although not actively discussed, the main underlying mechanisms are likely rural-urban migration, policy responses, or urban zoning in response to shocks that have been studied in previous literature (for some discussions see Cattaneo, Beine, et al., 2019; Ferreira, 2024). However, there are several concerns with the

[13] The number of times a hazard switches from non-treated to treated

[14] In literature, often clustering at larger levels is found (see for a discussion by Ortiz-Bobea, 2021 section 4.6). Given that part of the underlying mechanism, such as rural-urban migration, will be at the province level, this clustering might introduce correlation in the residuals. Clustering at the city level assumes errors are independent across cities, which may not fully capture this dependence. This is assessed in the baseline.

proposed methodology, notably, violation of the Stable Unit Treatment Value Assumption (SUTVA), omitted variable bias, reverse causality, spatial and temporal autocorrelation of the error term, and selection into treatment.

First and foremost, there is concern about violating the Stable Unit Treatment Value Assumption. This assumption states that the treatment assigned to a specific hazard is independent of the treatment status of other units. In the context of extreme events, this is often not the case. Potential reasons might be spatial spillovers, treatment anticipation or non-uniformly distributed treatment definitions. In this article, we consider anticipation to be limited. Although Early Warnings For All and other initiatives have substantially advanced access to information in recent years, their coverage remains unevenly distributed across our sample. Many countries still have only limited access to this information. As a result, anticipation might not be a major concern. Regarding treatment uniformity, we standardize treatments by event counts using mostly country-specific thresholds. This reduces concerns and heterogeneity of the treatment definition. However, spatial spillovers remain a valid threat to the SUTVA. Spillovers from nearby areas are possible (Chlouba, Mukim, and Zaveri, 2023), thus we include a test if treatments in catchment areas have more explanatory power. This will at least reduce concerns to some extent.

The omitted variable bias is present if time-varying factors that influence horizontal urban expansion are not captured in the model. Fixed effects help control for many omitted variables in the climate context. However, if shock variables come from the same distribution as climatic variables, omitted variable bias may still occur if key climatic variables are excluded (Hogan and Schlenker, 2024). Urban studies often examine the effects of rainfall and temperature on growth. However, air moisture, linked to heatwaves and urban disamenities, may also be important. Our use of pre-composed intensity measures (e.g., heatwaves or floods) helps reduce such bias, though some residual environmental factors remain. For example, powerful local weather systems, known as mesoscale convective systems, can lead to floods, storms, and hailstorms, an impact that is omitted in this article. From a socio-economic point of view, other predictors of urban growth at city level are excluded given that they might represent "bad controls", which means mediating the effect of the shocks (Dell, Jones, and Olken, 2014). It is not possible to fully account for all omitted variables, but we believe that with the set of fixed effects, the chosen controls as well as the inclusion of all events at once large parts of the omitted variable bias is limited.

Then, the impact of shocks on urban growth may be nonlinear. The current model assumes constant marginal effects (β_k), treating each additional event as having the

same impact — e.g. from zero to one heatwave is equal to three to four. This may oversimplify reality: if marginal effects vary with the number of events, unaccounted non-linearities remain in the error term as omitted variables. Ideally, we would employ a non-linear specification similar to that used by Burke, Hsiang, and Miguel (2015) by directly including intensities. However, this approach is more challenging in our case, as our model includes multiple main independent variables that may interact with each other when used to predict extreme events. Adding a non-linear function for rainfall extremes alone might be easy to implement. However, adding wet-bulb temperature, our measure for heatwaves in a non-linear way at the same time introduces significant complexity for interpretation. The relationship between these variables and their polynomials complicates isolating the impact of rainfall and wet-bulb temperature on observed effects. The chosen proxy (number of events per year) still retains at least an idea of the quantity of shocks, but allows a better identification of what shock drives the results. To at least partially address this limitation, our robustness checks include quadratic polynomial specifications and incorporate annual raw intensity measures.
Reverse causality is unlikely, as hazard shocks are lagged by one period. However, urban growth today may increase the likelihood of future heatwaves or flood (see e.g. Feng, Zhang, and Bourke, 2021; Tuholske et al., 2021). Increasing the risk of an event and causing an event are not the same thing. Floods and heatwaves depend on multiple factors beyond urbanized land, including city characteristics, geography, and climate. However, this introduces temporal autocorrelation of the error term. To address this, we apply spatial heteroskedastic-autocorrelated corrected (HAC) standard errors as a robustness check.
In addition, shock exposure varies spatially: droughts expand over large areas, while floods are often localized. Thus, our data likely captures urban flooding well but may underrepresent drought exposure, introducing spatial measurement error. To address this and account for the larger spatial footprint of some shocks, one of our extended heterogeneity analyses expands the definition of a shock from the city to its surrounding catchment area using geographic buffers.
A final concern is selection into treatment. In the baseline model, some cities in the control group are never exposed to any of the four shocks, which may have fundamentally shaped their growth trajectories. To assess whether this influences our baseline results, we apply the method proposed by Imai and Kim (2019), which constructs control groups with similar treatment histories and adjusts them based on relevant covariates. This approach enables us to assess the potential impact of selection bias on our findings.

4.4 Results

In the following sections, we first present our baseline estimates, followed by analyses of heterogeneity and robustness. We then introduce three key extensions. First, we disentangle the number of hazard events from the co-occurrence of other events, commonly referred to as "compound events". Second, extend the spatial scope beyond city boundaries to examine intensive margin effects by exploiting variation in hazard exposure across surrounding catchment areas. Third, we incorporate temporal dynamics through an event study that accounts for the history of treatment, allowing us to explore longer-term effects and address potential selection bias.

Table 4.1
Toy Equation

	Growth of Impervious Surfaces			
	(1) Flood	(2) Heatwave	(3) Drought	(4) Storm
$\sum$ Events in t-1	-0.0029*** [-0.004,-0.002]	-0.0016*** [-0.002,-0.001]	0.0011** [0.000,0.002]	0.0038* [0.000,0.008]
Adjusted R^2	0.223	0.223	0.223	0.223
Obs.	107,796	107,796	107,796	107,796
Average Urban Growth	.03	.03	.03	.03
Controls?	Yes	Yes	Yes	Yes
City-FE	Yes	Yes	Yes	Yes
Year-FE	Yes	Yes	Yes	Yes
Region-Year FE	No	No	No	No

Note: *** 0.01,** 0.05, * 0.1 p-val. Confidence intervals are at a 90% level. This table shows the results of Equation 4.1 excluding region-year FE and only adding hazards individually. Controls included are country employment in agriculture (in %) and the country's total GDP per capita. Regions are defined as subregions of Africa and are shown in Figure D.2.

Before presenting the results of Equation 4.1, Table 4.1 illustrates the common approach in the literature, which involves including hazards individually. This toy regression is similar to Equation 4.1 but includes only one hazard at a time and excludes region-year fixed effects. This shall allow us to illustrate the raw correlations in a two-way fixed-effects setting. Each flood leads to a reduction in urban growth by around 0.3pp (or 10% relative to the mean growth rate), heatwaves by 0.2pp (or 7% to mean growth), while droughts increase them by 0.1pp (+3% to mean growth) and storms of 0.4pp. (+13% to means). In the next sections, we are presenting how these individual coef-

ficients change once accounting for the mean level of the other hazards and will see that especially the heatwave becomes insignificant with droughts and storms losing magnitude and significance.

4.4.1 Main Results

We find that flood events significantly reduce the rate of urban expansion. Our estimates suggest reductions in urban growth of approximately 0.24 to 0.28 percentage points (pp) for each additional flood experienced, representing a reduction of around 8 to 9% relative to the city's average rate of growth of 3% (with a 3 to 9% reduction in the confidence intervals). This effect is robust to all tested specifications and is the only hazard that exhibits a clear pattern. Table 4.2 reports results from Equation 4.1, which examine the short-term effects of weather shocks on the intensive margin of additional shocks on urban expansion. We gradually increase the restrictiveness of our model from column (1) to (4). Column (1) shows a simple two-way fixed effects model, often found in the literature. Floods and heatwaves significantly reduce urban growth at the 1% level, while droughts and storms appear to increase it.

However, these positive effects do not hold when we account for spatial and temporal autocorrelation in column (2), using a 100 km spatial and 2-year temporal clustering window[15]. This suggests the importance of spatial dependencies, particularly for drought and storm events. We study this insight in more detail in Table 4.4.2. In column (3), we present our preferred specification, which is equivalent to Equation 4.1. We introduce region-year fixed effects and use clustering of standard errors at the city level. Here, the flood effect remains negative and significant, while the effects of droughts and storms become weaker. In contrast to Van der Borght et al. (2025) the effect of droughts is with 0.08pp on a 10% level substantially smaller than in literature, likely due to the inclusion of smaller cities in the sample. Finally, clustering at the country level in column (4) renders all coefficients but flood insignificant. This emphasizes that substantial heterogeneity exists within countries.

Heterogeneity

In the following paragraphs, we investigate whether the results vary when considering stratified samples from the entire city sample. All results are still based on Equation 4.1,

[15] Unfortunately, this clustering does not allow for more than two fixed effects, which is why it is applied to the two-way fixed effects model rather than the one with included regional fixed-effects.

Table 4.2
Baseline Regression

	Growth of Impervious Surfaces			
	(1)	(2)	(3)	(4)
∑ Flood Events in t-1	-0.0028***	-0.0028***	-0.0024***	-0.0024**
	[-0.004,-0.002]	[-0.004,-0.001]	[-0.003,-0.001]	[-0.004,-0.000]
∑ Heatwave Events in t-1	-0.0015***	-0.0015**	-0.0008	-0.0008
	[-0.002,-0.001]	[-0.003,-0.000]	[-0.002,0.000]	[-0.002,0.001]
∑ Drought Events in t-1	0.0011**	0.0011	0.0008*	0.0008
	[0.000,0.002]	[-0.000,0.002]	[0.000,0.002]	[-0.001,0.003]
∑ Storm Events in t-1	0.0038*	0.0038	0.0036*	0.0036
	[0.000,0.008]	[-0.000,0.008]	[0.000,0.007]	[-0.001,0.008]
Adjusted R^2	0.223		0.233	0.233
Obs.	107,796	107,797	107,796	107,796
Average Urban Growth	.03	.03	.03	.03
Controls?	Yes	Yes	Yes	Yes
City-FE	Yes	Yes	Yes	Yes
Year-FE	Yes	Yes	Yes	Yes
Region-Year FE	No	No	Yes	Yes
Clustered at	City	HAC	City	Country

Note: *** 0.01,** 0.05, * 0.1 p-val. Confidence intervals are at a 90% level. This table shows the baseline results of Equation 4.1. Controls included are country employment in agriculture (in %) and the country's total GDP per capita. Column (2) applies spatial HAC errors following Colella et al. (2019) with a distance cut-off of 100km and with 2 temporal lags. Regions are defined as subregions of Africa and are shown in Figure D.2.

but are divided into subsamples.

By Geographic Region Following the IPCC report (Seneviratne, Nicholls, et al., 2012), the different regions of Africa are differently exposed to extreme events, which motivates the division of the country into broader regions. The negative effects of floods are found to be strongly driven by Western African cities, as shown in Table 4.3. In these cities, each additional shock reduces growth by 0.4 pp. Droughts accelerate urban growth especially in Eastern Africa, where they increase urban growth by approximately 14% relative to the means. However, in western Africa, they tend to significantly reduce urban growth at the 10% level. All models include fixed effects for city and year. Two important takeaways can be concluded from Table 4.3: first, there is strong heterogeneity in Sub-Saharan Africa, likely because weather shocks can be spatially centered around a specific region, and second, the direction of the effects may also depend on the region. In Western, Southern and Central Africa floods, heatwaves and droughts are generally harmful to urban expansion, albeit mostly not significant,

Table 4.3
Heterogeneity by Region

	(1) Southern Africa	(2) Eastern Africa	(3) Central Africa	(4) Western Africa
∑ Flood Events in t-1	-0.0014	0.0009	-0.0017	-0.0046***
	[-0.003,0.000]	[-0.001,0.003]	[-0.007,0.004]	[-0.006,-0.003]
∑ Heatwave Events in t-1	-0.0032	0.0040**	-0.0028	-0.0010
	[-0.007,0.000]	[0.001,0.007]	[-0.006,0.000]	[-0.002,0.000]
∑ Drought Events in t-1	-0.0003	0.0044***	0.0000	-0.0015*
	[-0.001,0.001]	[0.003,0.006]	[-0.003,0.003]	[-0.003,-0.000]
∑ Storm Events in t-1	0.0042	0.0015	0.0000	0.0000
	[-0.001,0.009]	[-0.002,0.005]	[0.000,0.000]	[0.000,0.000]
Obs.	19,551	28,275	16,716	43,254
Average Urban Growth	.02	.03	.03	.03
Controls?	Yes	Yes	Yes	Yes
City-FE	Yes	Yes	Yes	Yes
Year-FE	Yes	Yes	Yes	Yes
Region-Year FE	No	No	No	No

Note: *** 0.01,** 0.05, * 0.1 p-val. Confidence intervals are at a 90% level. This table shows the baseline results of Equation 4.1. All samples include controls (a country's employment in agriculture (in %) and the country's total GDP per capita) and have city and year fixed effects. Regions are defined as subregions of Africa and are shown in Figure D.2 and countries are listed in Table D.1. The exposure by country is shown in Figure 4.3.

while the opposite is true in Eastern Africa. This reveals that the underlying mechanisms that might drive or hinder urban expansions are different. Regions may differ in their regional governance, their ability to respond to migration in the face of shocks, or their differences in urban absorptive capacity. Cities in East Africa might have potentially better absorptive capacity, or rural farmers tend to respond faster with rural-urban migration . In other parts of Africa, these channels might be less pronounced, or other factors such as rising conflict might be different responses to weather shocks keeping people from moving into cities. In Appendix D.4 we additionally assess if global climate phenomena that influence regional climate events might influence urban expansion. Specifically, we focus on the El Niño Southern Oscillation, which might amplify rainfall in some regions of Africa and increase dry conditions in others. Thus, as an additional test, the effects are analyzed if the year of the event is an El Niño year or not. The model and the data are described in the respective section. The findings show that effects are most pronounced in neutral years, where communities or policy-makers might anticipate less extreme events. There is regional variation in responses to events, depending on the year in the El Niño cycle. However, more temporally

disaggregated studies are needed to draw robust conclusions from these results.

By Urban Poverty The effects on the response of urban areas could depend on the ability of residents to cope with shock, especially in the context of poor households (Hallegatte, Vogt-Schilb, Bangalore, et al., 2017; Cattaneo, Beine, et al., 2019). Richer cities might have more capabilities to adjust to shocks, whereas poorer cities might respond rather than anticipate a weather shock. To this end, we use data on poverty and wealth measures at the village level by Lee and Braithwaite (2022) and compute the time-invariant wealth average of a city. The authors provide the International Wealth Index, which we use for this study. Based on the time-invariant measure at the city level, we construct four quantiles within each country, which show the least (richest) and most (poorest) deprived cities within each country in our sample. The results are shown in Figure 4.4(a). Flood has negative effects throughout the entire wealth distribution. A likely explanation is that flooding might lead to displacement and destruction within the city, independently of the city's wealth. These factors can cause the effects of growth to be negative. Droughts increase urban expansion only in the poorest cities; for all others, they appear to have no significant impact. A conjectured explanation is that rural-urban migrants might first settle in nearby agglomerations, as the cost of settling is potentially lower than migrating into richer, larger cities. Storms show a clear pattern, where richer cities seem to expand more after severe storms. This could be because storms are destructive and manifest more frequently near coastlines, where wealthier cities are often located. These factors prompt infrastructure investment after such events and lead to real estate redevelopment, which might drive urban expansion.

By Remoteness Usually, cities that are located in proximity to the main city of a country, the "metropole", might be better connected through trade and transportation networks. Households in remote cities might be poorer (Provenzano, 2024) and thus be more severely affected if help takes longer to arrive or their remoteness makes providing disaster aid less politically interesting (Hamza, Eriksson, and Staupe-Delgado, 2021). Thus, we computed the distance of each city to the metropole of the country and built quantiles of distance within each country. The only clear pattern is observed in Figure 4.4(b) for floods, where the strongest negative effects are found for the most remote cities. Combined with the previous result, it becomes clear that the largest loss in growth in response to floods is paid by remote, poorer cities. No effects are found for heatwaves. For storms, it is mostly cities close to metropolitan areas that experience increased growth. This is again in line with previous results, which show that

the richest cities tend to see the largest growth effects. For droughts, the results are more nuanced, where only cities at a medium distance see the positive effects. Overall, it seems that remoteness plays a role especially for flood events, where more remote cities see larger losses in growth.

By City Rank Van der Borght et al. (2025) find that droughts seem to especially increase growth in large economic centers. To test whether, in addition to wealth, the relative size of a city matters, we stratify our sample by creating a ranking of the top 3, 5, 10, and all other cities within a country[16]. One might imagine that large cities offer more opportunities for work and greater flexibility in relocating internally after an event. The results are shown in Figure 4.4(c). The results for floods are consistent with the previous sections. For heatwaves, we observe weak effects that indicate the largest cities in a country also experience the largest relative expansion after heatwave events. For drought, the contrary is the case. If large cities are affected compared in comparison to less affected ones, the effects of a drought reduce urban expansion. Storm has a sample too small to be evaluated here.

By Industry Type Lastly, we also test the proposed finding by Henderson, Storeygard, and Deichmann (2017) that conjectures that negative shocks of soil moisture contributed especially in manufacturing and export-oriented towns to increased urbanization. According to their argument, this is especially due to the possibility of manufacturing cities to absorb incoming labor from rural-urban migration. However, in all non-manufacturing towns, hardly any changes are observed. Similarly, Colmer (2021) finds that for India the possibility for firms to absorb incoming labor substantially reduced economic losses after rainfall and temperature shocks. In our context, manufacturing cities are defined as cities where more than 20% of total carbon emissions are derived from industrial production and processes[17]. This data is taken from the European Commission - JRC and International Energy Agency (IEA) (2023) and computed as sectoral averages by city-year. For floods and heatwaves, we find that cities that are classified as non-manufacturing see a minimal reduction in urban growth after each additional event in Figure 4.4(d). Droughts, similar to Henderson, Storeygard, and Deichmann (2017), have a positive effect especially in manufacturing cities. This is surprising in light of the results shown above. Previously, we found that espe-

[16] The ranking is build as average urban extent over the 20-year period within the country.

[17] The threshold has been set after assessing this metric for a random sample of 10 manufacturing cities taken from the replication package of Henderson, Storeygard, and Deichmann (2017). The share of industrial CO_2 fluctuated between 18% and 60%. The 20% is used as a lower bound estimate.

cially small cities experience expansions after droughts, while large cities observe the opposite. Manufacturing cities tend to be larger cities also, so one would expect negative, not positive, effects. This is likely driven by the change in sample. In the city-rank analysis, around 130 cities fall into the top 3 tier, while in this figure, one of every five cities fits the criteria. An additional robustness test, excluded here, revealed that the positive effects are driven by medium-sized cities with ranks between 10 and 50 that are sufficiently industrialized but not the primary cities of a country. For storms, we observe large confidence intervals, given the limited sample size. In contrast to Henderson, Storeygard, and Deichmann (2017), however, we include a larger sample of cities, covering a shorter but more recent period, and focus on a more recent classification of cities by manufacturing status. Nonetheless, despite different magnitudes, the direction is comparable. We do not test for the underlying mechanism that might explain why manufacturing cities see stronger growth responses. In the Indian context, short-term weather shocks (Colmer, 2021) and long-term climate shocks (Liu, Shamdasani, and Taraz, 2023) caused by temperature anomalies have led to labor adjustments and rural-urban migration in response. Here, it has been found that short-term weather shocks have led to reallocation from agricultural to non-agricultural employment, which is not observed in the long run. Similarly, as noted by Colmer (2021), similar effects may also be observed in an African context for other short-term weather shocks. A possible explanation might therefore be that drought induces more rural-urban migration, while floods cause temporary displacement (Cattaneo, Beine, et al., 2019). If migrants choose a destination, they might relocate to cities with larger labor market opportunities and thus settle in cities with more non-agricultural markets, thus accelerating growth in manufacturing cities. Floods or heatwaves might, especially in rural towns in Africa, trap the population due to losses of income, thus temporarily reducing urban growth in cities that do not strongly rely on manufacturing services. However, this finding is only indicative evidence. Further studies using firm-level data are needed to draw definitive conclusions.

In general, the results of this section reveal a large heterogeneity in the dimension of the impact of hazards on urban expansion in Africa. Western Africa experiences the most significant overall negative impacts, particularly from floods. These floods reduce urban expansion, especially for poorer, remote cities that tend to depend more on non-manufacturing industries. Droughts tend to drive urban expansion in the poorest and medium-distance cities. This is likely due to displacement-driven rural-urban migration. In general, this section highlights the importance of considering spatial and

economic context in examining the effects of urban expansion. Papers in the literature often pool cities only by continents. However, there may be significant heterogeneity in the impact response due to the characteristics of the city and region. This should be taken into account.

Robustness Checks

In order to ensure that our results are not driven by specifications or sample selection, a set of robustness checks is conducted.

Definition of Outcome In the literature, one of the most common annual data on urban expansion is the World Settlement Footprint (WSF) Evolution (WSFevo) from 1985 to 2015 provided by Marconcini et al. (2020) (see e.g. studies by Chlouba, Mukim, and Zaveri, 2023; Rentschler, Avner, et al., 2023). The method of obtaining WSFevo slightly varies from the GISA data used in this study. Column (1) of Table 4.4 if the results hold using the WSFevo data instead. The results are comparable to the baseline, except for droughts, which become insignificant. WSFevo and the used GISA data have the shortcoming that they always increase, which might overlook urban destruction. To this end, column (2) measures the urban growth in a variable that can regress: low-resolution radar backscattering. In this method, radar signals are sent and received from satellites, and the stronger the signal received, the greater the potential build-up. For this purpose, we use a homogenized time series of Tao et al. (2023). After computing growth rates based on average city-year signal[18], Column (2) of Table 4.4 shows that effects that allow destruction in urban land do not fundamentally change the baseline results. In contrast, it makes them even stronger. In general, it appears that the measure of urban expansion is robust against other measures.

Definition of Growth Two closely related papers by Chlouba, Mukim, and Zaveri (2023) and Van der Borght et al. (2025) use different definitions of growth. To avoid the definition of growth driving the results, the two definitions of growth are tested. Column (3) of Table 4.4 tests the effects using an increment plus a constant factor, and Column (4) relies on longer-term growth, assuming that extreme events, in the case of

[18] Note that backscatter might also increase due to high-rise urban build-up, a phenomenon known as "double bouncing." It can also increase through high surface backscatter from flooded areas. Most cities in Africa have a fairly low share of high-rise buildings. In addition, annual backscatter statistics help smooth out noise from flooded areas. Therefore, this is likely a minor concern in this robustness test.

Table 4.4
Robustness Tests of Baseline Model

	Alt. Outcomes		Alt. Growth Specifications		Alt. Structural Form
	(1) WSFevo	(2) Radar Signal	(3) log(Δ Increment+1)	(4) 2-year Growth	(5) Standard
∑ Flood Events in t-1	-0.002** [-0.003,-0.000]	-0.003*** [-0.003,-0.002]	-0.083*** [-0.121,-0.045]	-0.007*** [-0.011,-0.003]	-0.002*** [-0.003,-0.001]
∑ Heatwave Events in t-1	-0.001 [-0.001,0.000]	-0.000** [-0.001,-0.000]	-0.028 [-0.066,0.010]	0.006* [0.000,0.011]	-0.001* [-0.002,-0.000]
∑ Drought Events in t-1	0.000 [-0.001,0.001]	0.008*** [0.008,0.009]	0.004 [-0.026,0.034]	0.003 [-0.000,0.006]	0.001 [-0.000,0.002]
∑ Storm Events in t-1	-0.001 [-0.004,0.002]	-0.004*** [-0.007,-0.002]	0.117** [0.021,0.213]	-0.005 [-0.014,0.005]	0.004* [0.000,0.007]
Build-Up (in km^2) in t-1					-0.000*** [-0.001,-0.000]
Obs.	84,977	100,986	108,907	96,520	107,796
Average Urban Growth	.03	.03	9.4	.07	.03

Note: *** 0.01,** 0.05, * 0.1 p-val. Confidence intervals are at a 90% level. This table shows the baseline results of Equation 4.1. All samples include controls (a country's employment in agriculture (in %) and the country's total GDP per capita) and have city and year fixed effects. Regions are defined as subregions of Africa and are shown in Figure D.2. Column (1) includes data from the German Aerospace Center (DLR) (2024) and Marconcini et al. (2020) measuring build-up growth until 2015. Column (2) utilizes low-resolution backscattering radar signals from unified time series from Tao et al. (2023) as an alternative measure of construction growth. Columns (3) and (4) change the outcome transformation, and Column (5) adds a lagged term into the model.

Van der Borght et al. (2025) droughts, take place to affect urbanization. The change does not fundamentally change the result.

Growth Regression Framework To align with the growth literature following, e.g., Mankiw, Romer, and Weil (1992), Column (5) of Table 4.4 includes a lagged level of the dependent variable. This specification yields results that are statistically more significant, suggesting that the baseline results are robust to this specification change. The model tested might be subject to the so-called "Nickell bias" (Nickell, 1981) referring to concerns that in short panel models when lagged dependent variables are included as regressors in a fixed-effects model. Here, the demeaning due to individual fixed effects induces a correlation between the transformed lagged dependent variables and the transformed error term, which biases the estimates and renders them inconsistent. To avoid drawing incorrect conclusions, we also apply a bias-corrected model by (Breitung, Kripfganz, and Hayakawa, 2022), which yields comparable results.

Sample Dependence Most articles in the literature apply some sort of threshold to only include a fraction of the sample. We test in Table D.2 whether the results change if we winsorize the dependent variable by applying a 5% threshold in Column (1). The results become weaker for floods and stronger for storms, but the direction of the result remains the same. Next, we progressively excluded cities from columns (2) to (4) that had populations below a certain threshold by the end of 2019. The results for floods are persistent around 0.2-0.3pp excluding cities with urban population of less than 5 000, 10 000[19] and 50 000 people. As only around 20% of the cities in the sample have more than 50 000 people in 2019, this test mainly excludes large parts of the sample. Including only small cities with less than 50 000 people or less in 2019 we have similar effects in magnitude, but droughts have a more significant positive effect of increasing urban expansion by 0.1pp. This shows that part of the positive baseline effects of droughts are driven by small cities, which was already discussed in Figure 4.4(c). In general, the main claim is that floods, especially, are causing a reduction in urban expansion. This claim is strongly robust against sample dependence and outliers in growth. To also test if one country drives the results, we perform a leave-one-out exercise as shown in Figure D.1. The results for floods remain significant regardless of which country is dropped. For heatwaves, it seems that Nigeria and the Ivory Coast drive the results

[19] Note that one criteria for cities to be included in Africapolis (OECD and SWAC (2020)) is that more than 10 000 people live in a city. The data used here to compute urban population is different to the one of Africapolis, thus we might find still urban agglomerations with less than 10 000 people. Some cities may also have lost population over the years

more significantly. For drought, excluding Nigeria, the results are more positive and significant. For storm, since the largest treated sample is in South Africa, excluding it makes the storm insignificant. This reveals that the results for floods are stable across countries. For heatwaves and droughts, Nigeria partially drives the results, which, given that it accounts for 20% of the sample, is not entirely unlikely.

Measure of Shock The main specification of the article sums all events of the four hazards in a year over an entire year. Two critiques should be tested: first, an earlier shock in the year might have a greater impact on urban growth. Second, the intensity (duration or physical intensity) might also matter. To address the first concern, in column (3) of Table D.3, we compute a time-decay weight[20] that gives a stronger weight to earlier occurring events. The events become more significant and marginally larger, especially for floods, heatwaves and droughts. The weighting of the shock has an influence, but the direction of the results remains. To address concerns about intensity, column (5) includes, instead of the number of events, the number of months in a year that have conditions exceeding a threshold. This captures instead of the frequency of new onsets, the duration of onsets. Only the effects for floods are significant and smaller. For other hazards that are insignificant in terms of duration, the impacts may be more centered around the onset rather than the duration of an event. Lastly, instead of including event counts, we add the raw intensity measures aggregated at the yearly level, which may oversimplify the effects. For floods, we include the maximum DFO-Index in a city within a year, for heatwave the highest recorded 5-day block maximum in terms of Wet Bulb Temperature, for drought the lowest recorded SPEI-3 and for storm the highest daily wind gusts in meters per second[21]. The effects are around 0.09pp smaller as in the baseline in Column (1). Heatwaves become significant, probably because not accounting for historical deviations a large fraction of the sample sees already fairly high levels of temperature (most of Western Africa). An increase of one unit in SPEI-03 reduces urban growth, while the storm is still significant. Although effects on floods and storms remain, the direction of effects for droughts and heatwaves is reversed. A likely explanation is either that intensities affect urban expansion in a non-linear man-

[20] The weights are computed by month of the event as:

$$\frac{13 - (month)}{12}$$

[21] Note that in the conversion for the events, heatwaves also include deviations from historic means of months and droughts also include SPEI-12 beside SPEI-03

ner or that measurement errors occur. This is given because the Min-Max values of the intensities can be fairly noisy. We must be careful in interpreting these intensities as causal for two reasons. First, we underestimate the actual exposure, given that we only consider the maximum or minimum values over a year. This means that we only capture the most extreme events, underpredicting weaker, yet still impactful events. Additionally, especially for heatwaves and droughts, we do not account for historical deviations (for droughts for longer-term droughts). Overall, this section has shown that the results are fairly robust to variations in the definition and weighting of the shock variable.

Non-linearities Our model might have an omitted variable bias if non-linearities are not accounted for. Building on the results from the previous robustness test, we also add squared terms of all variables into all the described models. Column (2) of Table D.3 tests adding squared terms to our main specification. Floods, become are sitll significant but now only on a 10% level, no other coefficient is significant. Column (4) applies the same to time-decay-weighted shocks. Here, we find non-linearity in floods where having more than a single event tends to reverse the growth effects. No other hazard is significant. The same results are found when using the number of months squared in Column (6). In Column (8), accounting for squared non-linearities in intensities turns the flood entirely insignificant. However, for raw intensities, heatwaves and storms change the direction of the results, accounting for their squared terms. This highlights that for flood intensity, the linear and non-linear terms might offset each other (due to strong correlation), making both insignificant, or there is no relationship across intensity levels[22]. Heatwaves and storms seem to have non-linear dependencies. The key takeaway from this robustness analysis is that our baseline measure, annual event counts, is relatively robust to alternative shock definitions and the inclusion of non-linear terms. However, we find some indication that the economic impacts of floods are both nonlinear and temporally sensitive, particularly when considering earlier events in a given year. This illustrates that aggregating shocks into an annual, equal-weighted count of floods may overlook important variations related to the timing of events. Additionally, this approach may overlook variations in the frequency of events within a given year. Given that our empirical strategy relies on a panel model at the yearly level and we assume that shocks take some time to influence urban ex-

[22] A likely explanation can be found looking at the range of considered intensities. The values range from 0 (no flood) to 2 (extreme flood). Most flooded years have an intensity of one. This causes the squared term to be collinear to the flood event and being dropped. This might explain at least in part the insignificant results.

pansion, we continue to use annual event counts in our main specification. However, we acknowledge that this approach may overlook some time-sensitive non-linear dynamics, especially for floods.

Significance of Floods What is striking in the previous results is the robustness of flood events. Thus, additional robustness tests for flood events are performed in Table D.4 to test more or less restrictive flood definitions. Neither including more common floods nor excluding cities with little or no flood risk changes the results. To avoid assigning floods to non-treated cities, we furthermore build more restrictive measures of flood. These measures consider whether flood records co-occurred with extreme rainfall shocks (pluvial floods) or with extreme river streamflow discharges (riverine floods). Both effects are significant and negative, but pluvial floods have notably stronger impacts.

4.4.2 Beyond the Baseline: Three Extensions

The previous section focused on the baseline equation in Equation 4.1. The following sections introduce three extensions, notably the definition of a shock, before turning to a discussion and interpretation of the results. First, the shock is decomposed into cumulative hazard shocks and interconnected shocks, referred to as "compound events." Second, the analysis examines whether the results change when shifting from an intensive margin at the city level to an extensive margin at a broader spatial scale. Third, the temporal dynamics of the proposed shock are explored.

Accounting for the Shock Complexity

Following the work of, for instance, Zscheischler, Martius, et al. (2020), accounting only for four distinct hazards might oversimplify the true nature of the effect of shocks. This part will study what happens if shocks of different types co-occur, e.g. a flood following a heatwave. These combinations might exacerbate the impact of shocks and thus lead to significantly different results. We categorize shocks into single-hazard and multi-hazard events, based on whether one or more hazards occurred within a month or a given time window. We define "compound" events as occurrences where multiple hazard types (e.g. flood and storm) are observed within a single month or a 2-month span. Furthermore, we also test for the two-month window the order in which both event classes took place. More information on the creation of the shock dummy and

Table 4.5
Extension: Accounting for Compound Events

	No Compound Events (CE)		Not ordered		Ordered
	(1)	(2)	(3)	(4)	(5)
∑ Flood Events ex. CE	-0.002*** [-0.003,-0.001]	-0.002*** [-0.003,-0.001]	-0.002*** [-0.003,-0.001]	-0.002*** [-0.003,-0.001]	-0.002*** [-0.003,-0.001]
∑ Heat Events ex. CE	-0.001* [-0.002,-0.000]	-0.001* [-0.002,-0.000]	-0.001* [-0.002,-0.000]	-0.001* [-0.002,-0.000]	-0.001 [-0.002,0.000]
∑ Drought Events ex. CE	0.001 [-0.000,0.002]	0.001 [-0.000,0.002]	0.001 [-0.000,0.002]	0.001 [-0.000,0.002]	0.001 [-0.000,0.002]
∑ Storm Events ex. CE	0.003 [-0.000,0.007]	0.003 [-0.000,0.007]	0.004 [-0.000,0.007]	0.003 [-0.001,0.007]	0.003 [-0.001,0.007]
Multivariate Compound Events					
Flood × Heat			-0.002 [-0.006,0.002]	-0.004* [-0.007,-0.000]	
Flood × Drought			-0.007* [-0.014,-0.000]	0.001 [-0.003,0.006]	
Flood × Storm			0.006* [0.000,0.011]	0.006** [0.002,0.010]	
Heat × Drought			0.004 [-0.002,0.010]	0.002 [-0.003,0.008]	
Drought × Storm			0.001 [-0.001,0.003]	0.002 [-0.001,0.004]	
Heat × Storm			0.015*** [0.013,0.017]	0.015*** [0.013,0.017]	
Flood → Heatwave					0.004 [-0.000,0.009]
Heatwave → Drought					0.001 [-0.006,0.008]
Flood → Storm					0.010** [0.003,0.018]
Storm → Drought					0.003 [-0.005,0.012]
Drought → Flood					0.002 [-0.002,0.006]
Drought → Heatwave					0.001 [-0.005,0.007]
Drought → Storm					-0.001 [-0.006,0.003]
Flood → Drought					-0.009* [-0.017,-0.000]
Heatwave → Storm					0.006 [-0.005,0.016]
Heatwave → Flood					-0.007*** [-0.009,-0.005]
Storm → Flood					0.006** [0.002,0.010]
Obs.	107,796	107,796	107,796	107,796	107,796
City-FE	Yes	Yes	Yes	Yes	Yes
Year-FE	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes
Clustered at	City	City	City	City	City
Window to Define CE	1-month	2-month	1-month	2-month	2-month

Note: *** 0.01,** 0.05, * 0.1 p-val. The table illustrates the impact of compound events on physical urban expansion, based on an adjusted Equation 4.1. More information on the creation of the variables is given in D.5.

the updated regression model can be found in D.5. Then the baseline model is re-run now including these compound shocks. The results can be seen in Table 4.5. Columns (1) and (2) show the results when the baseline model's shocks omit months with compounded shocks in the same month (column (1)) or within a 2-month period (column (2)). Excluding compound events, the coefficients are about the same magnitude as in the baseline. Next, adding compound dummies in columns (3) and (4), the effects remain stable. The compound interactions of flood-storm and heat-storm are the most significant and both positive. The combination of droughts and floods shows a significant reduction of urban expansion by 0.7pp, almost triple the effect than for floods alone. This effect is only found for compound events within the same month. Flood-storm and heat-storm pairs need to be interpreted with caution given that only around 10 cities have seen these hazard pairs[23]. Finally, in Column (5), accounting for the order of the sequence, again, the effects of individual hazards remain. The most significant effects are observed for floods that follow heatwaves. Each additional sequence of that order occurring in a year reduces urban expansion by 0.7pp. This reduction amounts to around 23% reduction from the mean urban growth rate. These types of compound events are the most common found interaction and have been shown in literature to exist (Sauter et al., 2023). While we would have expected to also find significant effects for droughts before floods, notably of dry soil increasing the flooding impact, the effects are insignificant. In contrast, floods before droughts show a weak yet still significant negative effect. Seemingly paradoxical, floods can precede droughts for two reasons: first, floods are local events while droughts are global events. We might have a city experiencing a flood, and in the following month the catchment of a city experiencing a drought. Second, a flood does not exclude the possibility of a drought given that rapid flash floods might run-off without soaking into the soil. Lastly, again with a small sample, storm and floods and floods and storms likely capture tropical storms that bring storm surges in coastal areas. Here effects are positive and large with 1pp increase in expansion. The main takeaway from this extension is that, while the general direction of the results holds, compound pairs, especially in the case of heatwave-flood interactions, may show a significantly larger response than either floods or heatwaves alone.

[23] See Appendix D.3

Accounting for Spatial Scope

The second extension is related to the spatial aggregation of the shock. In Equation 4.1, shocks are computed as an intensive margin ("What is the homogeneous short-term effect of an additional event of type k?"). This is especially useful for indirectly examining local drivers of urban expansion, such as reconstruction, zoning policies, or city-specific responses to local events (floods and heatwaves). For events, notably droughts, covering larger spatial environments, it may be worth studying catchment area dynamics rather than cities themselves, as done by Van der Borght et al. (2025) (country-level shocks) or Chlouba, Mukim, and Zaveri (2023) (shocks in neighboring 50 km cells). To account for this dynamic, the analysis of this article is repeated by computing the extensive margin ("What fraction of the surrounding land around a city is affected?"). For this purpose, buffers of 50, 100, and 150 km are drawn around each city and intersected with a 10x10 km grid. As the shock is now studied not categorically but in the affected area, the shocks are also recomputed at the city level. More information on how buffers and shocks are defined Appendix D.6. Buffers have the limitation that they do not account for the transportation infrastructure within. Thus, as robustness, we also include daily commuting zone data by Data for Good at Meta (2023).

The models presented in Table 4.9 follow Equation 4.5. The model used is similar to the original, but it estimates what would happen if every spatial unit (cell) within the catchment of a city were to experience an additional event. Thus, it now measures the average aggregate effect of a shock with widespread exposure. We control for the population living within that buffer to determine the exposure of the population within that buffer. Column (1) shows that with the new definition of the shock that defines shocks at the city level, suddenly all coefficients become significant, similar to the model without region-year fixed effects in Table 4.2. Floods are most significant, but, more importantly, droughts are also significant with 0.1pp, increasing average urban growth by around 3.6%. Columns (2) to (4) show that the effects become stronger for floods, droughts and storm. As robustness, considering commuting zones instead, all but droughts are significant. It remains unclear why droughts are insignificant. However, a conjectured hypothesis is that, given that commuting zones might cross country borders, underlying migration mechanisms through droughts might predominantly take place within countries. One needs to be careful in the coefficient interpretation, as the average level of shocks slightly vary one needs to be careful to directly compare shocks in different buffers. Although the average shock for each hazard re-

mains relatively stable across buffer sizes, the coefficients controlling for population in each buffer increase in magnitude. This pattern suggests that larger coefficients are likely due to the broader spatial impact of the shock rather than differences in shock means across buffers. Thus, part of the observed impact on urban expansion in the baseline area may be driven by dynamics occurring in the city's broader catchment area. These dynamics may involve hidden mechanisms such as rural-urban migration or the avoidance of flooded cities as a destination.

Accounting for Temporal Dynamics

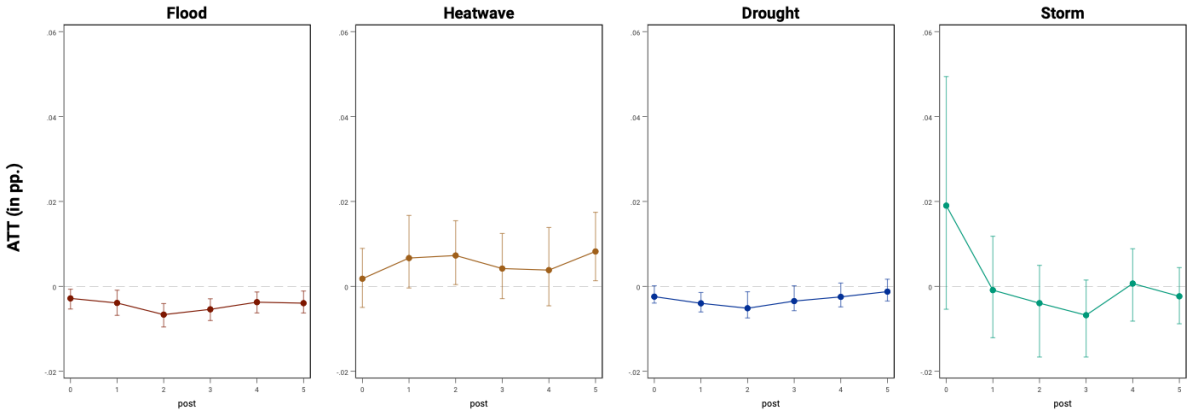

Fig. 4.2. Extension: Temporal Dynamics of the Weather Shocks

Note: The four panels show the average treatment effect of the treated (ATT) following the methodology by Imai and Kim (2019). The estimates use covariate propensity score matching to build valid counterfactuals and are matched on the eight previous treatment periods. Confidence intervals shown represent the 10% confidence intervals.

Third, and finally, all the models studied so far have studied the short-term impact of events on urban growth. However, researchers might be interested in the longer-term dynamics of shocks. In addition, it might be worthwhile to consider the history of treatment of a city to avoid selection bias of the results. To this end, we apply a recent estimator by Imai and Kim (2019). This estimator follows a three-step procedure. First, for each treated city by a hazard, a set of cities is created that share a similar treatment history over a specified pre-treatment period. In our case, for each treated city, we match five comparable cities with similar treatment histories over the last 8 years. Second, the comparable cities are matched to have a better balance between the treated city and the control set. In our case, we use the last two years of all other shocks than the one under consideration. Additionally, the total urban size in square km in 1990

and the average urban growth over the last four years are used to create a better control group. Finally, a difference-in-difference estimate is computed for post-treatment periods. More details are specified in Appendix D.7. The results of this exercise are shown in Figure 4.2. Accounting for the treatment history, floods not only lead in the year after, but over the full studied five post-treatment years to a negative effect on urban expansion. The negative effects increase from 0.2pp in the treatment year up to 0.6pp reduction three years after the event, all significant at a 10% level. Heatwaves show a positive, yet barely significant effect. Surprisingly, accounting for the treatment history of droughts, experiencing a drought decreases urban expansion by around 0.3pp in the year after a drought event up to two years. Then it turns insignificant. This is in contrast to the baseline results, where we document positive effects. That highlights that, while in the short run, droughts might have some positive effects. Especially for this hazard type, dynamic treatment matters, which simple growth models might miss. Thus, cumulative drought shocks over time must be taken into account. Lastly, for storms, no significant effects are found. The results are robust against different matching methods and placebo test to determine if pre-treatments exist. The key takeaway from this analysis is that the found effects for floods indeed are significant and might persist even longer than shown in the baseline. For droughts, it might be important to consider the dynamic treatments of these events. Taking into account the history of drought, the effects become negative.

4.4.3 Interpretation and Discussion

This section evaluates whether the changes in urban expansion have meaningful implications in the real world. Floods cause damage and potential loss of life. However, they may also prolong urban expansion. This creates an opportunity for communities to adapt by implementing non-structural measures, such as stricter zoning laws or incentivizing relocation. To understand the internal impacts of these events, we provide descriptive insights in Table 4.6 on how urban expansion dynamics are connected to population dynamics. More information on the data sources used and their transformations is shown in Appendix D.8. The baseline models in Table 4.2 demonstrate a robust decline in short-term urban growth rates following floods. However, this appears to coincide with an increase in population, approximately doubling the growth rate in Column (1). Increasing urban population in the context of climate have already been documented in literature (Castells-Quintana, Krause, and McDermott, 2021). The combination of slower build-up growth but strong population growth appears to in-

crease population density by up to 30 persons more per square km in Column (2). This effect is only found for floods. Furthermore, we observe a 64% increase to the means in population growth (or 3.2pp) in strongly informal settlements and almost 80% (or 4pp) increase in weakly informal areas of cities[24]. To our knowledge, this is a novel insight to the limited literature on how hazards influence informal growth. For a Colombian context, it was shown that informal settlements reduce in growth after extreme events (Camacho, Aryal, and Rajabifard, 2025). Surprisingly, Column (5) shows that population growth in previously flooded areas (or at least areas at high risk) almost doubles in the year following a flood. This might suggest people returning after a temporary displacement potentially, either due to lower housing costs or lack of regulation or enforcement of zoning laws or the housing market in these areas. This is in line with literature finding long-term growth of build-up in flood zones (Rentschler, Avner, et al., 2023). This effect is also, albeit weaker, found for droughts, where each additional drought event increases population in flood-risk zones by 22% to the mean growth (0.9pp). This result is also found with a different magnitude by Van der Borght et al. (2025). Finally, weather shocks might also influence violence and social unrest (Castells-Quintana, Lopez-Uribe, and McDermott, 2022; Eberle, Rohner, and Thoenig, 2025; Damette and Goutte, 2023). Column (6) tests the effects on urban unrest and violent protests. Here, instead of growth, we focus on the log of conflict events reported in the ACLED database (Raleigh, Kishi, and Linke, 2023). Floods and droughts significantly increase the number of violent unrests that occur within urban boundaries in the year after an event. This is in line with findings by Castells-Quintana, Lopez-Uribe, and McDermott (2022). There is indicative evidence that reduced urban growth does not necessarily translate into improvements in urban welfare after extreme events. Quite the contrary is true: especially floods, but also partially droughts, lead to increased population in informal neighborhoods and high flood-risk areas. They also increase the propensity for urban social unrest.

Combining all these dimensions allows us to place this article in the literature. Floods significantly reduce urban expansion, especially in more remote rural towns. This is similar to the findings on nightlights after flood impacts by Gandhi et al. (2022). Magontier and Martinez (2023) find for Spanish cities a vertical expansion in build-up, which is not tested in this paper. They also see increases in population in the aftermath

[24] The data is taken from the Million Neighborhoods project (Bettencourt and Marchio, 2023) and measures informal areas as gradient within neighborhoods that are badly connected to urban infrastructure.

Table 4.6
Within-City Impacts of Weather Shocks

	Population		Population Growth in ...			Other
	(1) Population Growth	(2) Population Density	(3) Informal Settlements (strong)	(4) Informal Settlements (weak)	(5) High-Risk Flood Areas	(6) Log(Conflict Count)
∑ Flood Events in t-1	0.042*** [0.035,0.049]	27.287*** [16.050,38.524]	0.032*** [0.024,0.041]	0.040*** [0.033,0.048]	0.042*** [0.035,0.050]	0.058* [0.009,0.107]
∑ Heatwave Events in t-1	0.004 [-0.001,0.008]	-0.976 [-12.753,10.801]	0.008* [0.001,0.015]	0.003 [-0.002,0.008]	0.003 [-0.002,0.008]	-0.051 [-0.108,0.005]
∑ Drought Events in t-1	0.006 [-0.001,0.012]	4.967 [-2.400,12.333]	-0.000 [-0.009,0.009]	0.005 [-0.001,0.012]	0.009** [0.002,0.016]	0.047** [0.011,0.084]
∑ Storm Events in t-1	-0.031 [-0.094,0.033]	12.562 [-23.472,48.597]	-0.060 [-0.170,0.051]	-0.027 [-0.085,0.032]	0.009 [-0.030,0.048]	-0.191 [-0.415,0.034]
Obs.	96,885	108,736	64,701	96,402	94,366	10,979
Average Growth of Y	.04	899.41	.05	.05	.04	1.1

Note: *** 0.01,** 0.05, * 0.1 p-val. Confidence intervals are at a 90% level. This table presents the baseline results of Equation 4.1 for different outcomes at the city level. Columns (1) and (2) show population growth and density based on estimates by Liu, Cao, et al. (2024a). Columns (3) and (4) estimate population within informal settlements with complexity index >10 (column (3)) and >6 (column (4)) from Bettencourt and Marchio (2023). Column (5) shows population in high-risk flood zone with inundation depths >50cm in 100-year floods following Rentschler, Salhab, and Jafino (2022). Column (6) uses conflict data from the ACLED database (Raleigh, Kishi, and Linke, 2023). Population data is interpolated as described in D.8.

of flood events, similar to our findings. Yet additionally, we find that they are amplified when co-occurring with rainfall extremes and following heatwave events. This is especially a concern for Western Africa. Heatwaves themselves have very mixed effects, and it is not clear if they contribute or not to urban expansion. We are not aware of studies studying the impact of heatwaves onto urban expansion as most studies focus on the inverse direction. Part of the results regarding droughts align with previous literature. Similarly to Chlouba, Mukim, and Zaveri (2023) and Van der Borght et al. (2025), we find coefficients that indicate that droughts accelerate urban expansion, especially taking into account catchment areas. There are, however, two major differences: first, as we account also for small cities, most positive growth effects are recorded in small cities, not in the primary cities of the country. This may be due to differences in sample sizes and varying definitions of the drought shock. Second, adjusting for the history of droughts, the effects become negative, highlighting that the longer-term effects of droughts might not be as positive, which has also been found by Nordfors (2024).

4.5 Conclusion

This paper aims to answer the question what the effects of four of the most important weather shocks in Sub-Saharan Africa are on urban expansion. In contrast to previous research, this study adopts a dedicated multi-hazard perspective, given the increased likelihood of co-occurring events in future decades. Using high-resolution data on urban development in combination with city-wide hazard exposure data for flood, heatwaves, droughts and storms, we reveal a complex pattern of the impact of extreme events across space, time and hazard type. Each additional flood reduces urban growth in the range of 3-9% relative to the average growth rate, especially in poorer, more remote towns, leading to increased population pressure in informal settlements and, surprisingly, parts of cities with high flooding risk. These effects are notably found in Western Africa. A novel finding is that these effects are fundamentally amplified when floods follow heatwave events, combinations that have been studied in atmospheric science before (Sauter et al., 2023). Similarly to previous studies, we find that drought increases urban expansion, but contrary to previous work (Van der Borght et al., 2025) mainly in smaller cities. For drought impacts, spatial scale matters. Although droughts may be harmful to the nearby agricultural land in urban areas, the more catchment areas affected, the stronger the effects on urban expansion. We discuss, however, that

these findings need to be regarded with care, as we reveal that accounting for the history droughts in cities the effects turn negative. This highlights that while in the short term some positive effects might be seen, accounting for longer-term effects, droughts tend to be more harmful to urban expansion. For heatwave, as expected, barely any significant effects are found, if so they tend to be rather negative. Extreme wind storms are geographically more centered leading to small treatment samples and large coefficients. Interpreted carefully, they seem to especially contribute to urban expansion in richer larger cities in southern Africa.

The paper is grounded in several assumptions whose implications should be considered when interpreting the results. A main limitation is the way how hazards are included. Most models assume a linear parametric short-term impact of shocks onto outcomes, which following the literature by Hogan and Schlenker (2024) might overlook important non-linearities. While tested for simple quadratic terms, the usage of intensity thresholds to define extreme events might omit this non-linearity. This is a significant limitation, as also points marginally below a threshold might be important and interact non-linearly with the model. Another limitation is that we only indirectly consider mechanisms such as rural-urban migration and within-city dynamics. A more rigorous micro-founded analysis is needed to substantiate the points made. These limitations provide several directions for future research, particularly non-linear modelling of events and spatial mechanisms such as migration.

The findings of this study have strong policy implications. Although adaptation action in Sub-Saharan Africa is currently focused on metropolitan areas, where a large fraction of the population is clustered, the findings suggest that small flood-prone cities should be more in focus for climate action. These cities tend to bear a large part of the burden of extreme events, with increasing informality and higher conflict potential. Furthermore, the analysis highlights that policymakers should refrain from conducting research exclusively on a whole-continent scale or along a single hazard dimension. The underpinning climate dynamics, such as teleconnections, might cause fundamentally different effects across subregions. Likewise, weather shocks do not occur in in isolation. Their impact also depends on the past treatment history of a hazard and their co-occurrence with other events, which may interact in non-linear ways. Ignoring such dynamics might mislead policy guidance.

This study is far from exhaustive. Beyond just adding non-linear models and spatial mechanisms, more research is needed to better study the individual and firm response on a large scale to extreme events, especially those that are expected to see a substantial

rise and interaction with others in the years to come. Future research should adopt a more inclusive approach that incorporates multiple types of hazards simultaneously. It should also connect with high-frequency microlevel estimates of response and damage reported to get a better overall picture of the complex heterogeneous response to climatic shock. Given the current climatic predictions that the 2.5-degree global warming goal is likely to be missed, it is more important than ever to start implementing adaptation policies in regions and urban agglomerations with less resilient infrastructure to safeguard human lives.

4.6 Figures & Tables

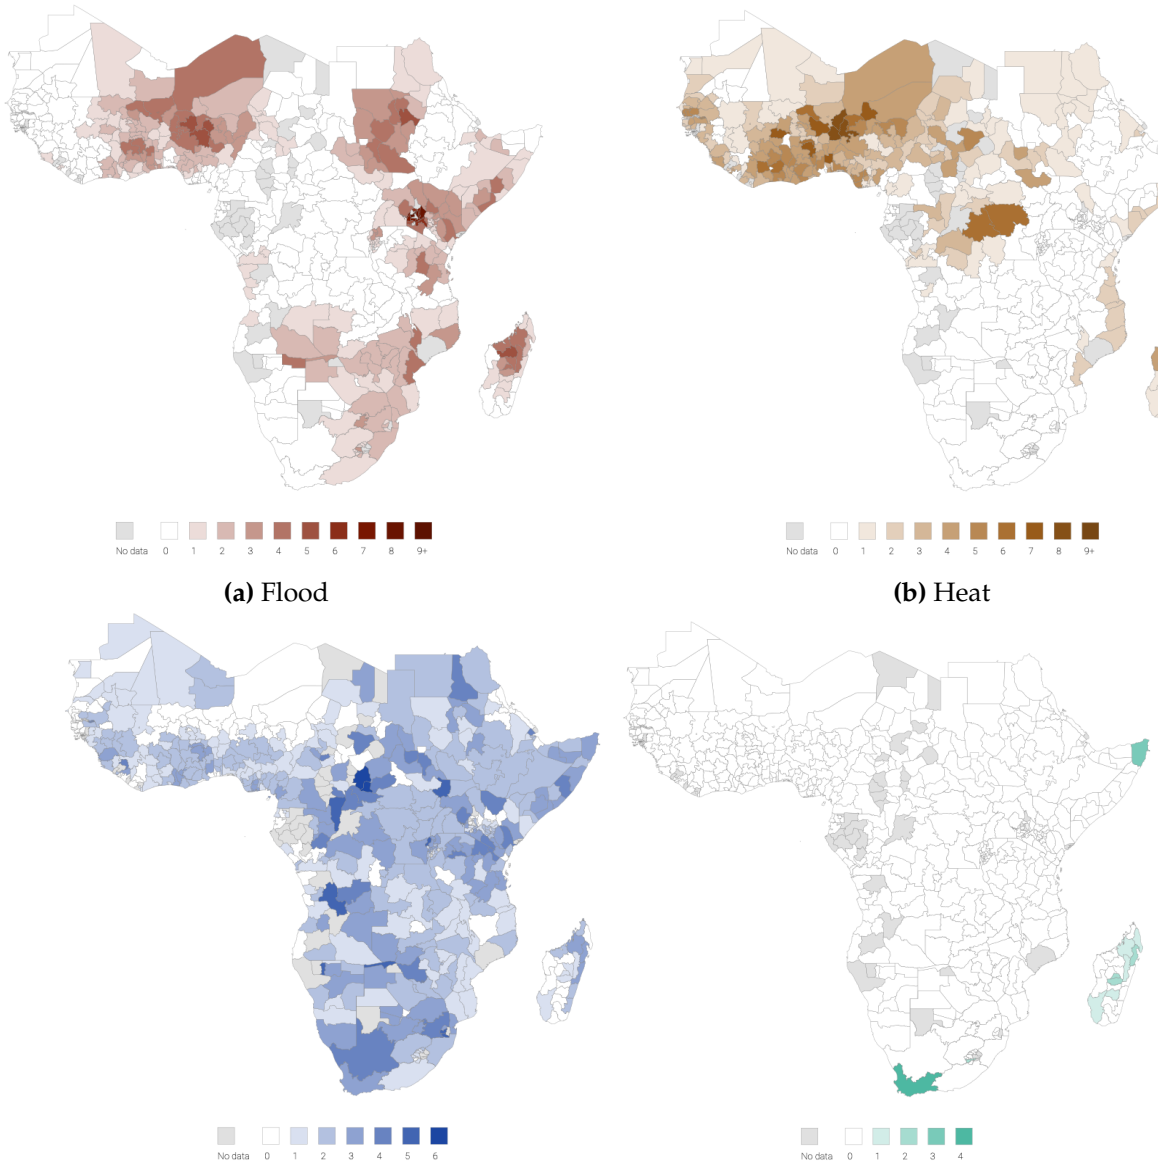

(a) Flood **(b)** Heat

(c) Drought **(d)** Storm

Fig. 4.3. Spatial Distribution of the Hazards

Note: The four figures show on a province-level (ADM 1) the average number of total shocks cities have seen between 2000 and 2019 across Sub-Saharan Africa. The data is derived from Reinhardt (2024a) and provinces are defined as in GeoBoundaries (Runfola et al. (2020)).

Fig. 4.4. Heterogeneity Analysis - Part 1/2

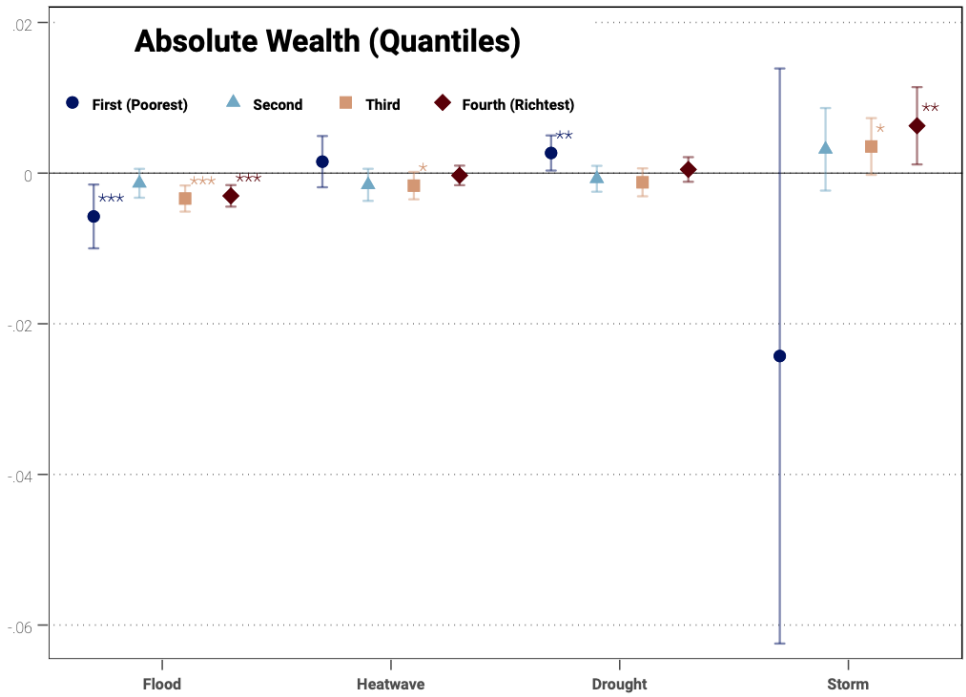

(a) By Absolute Wealth

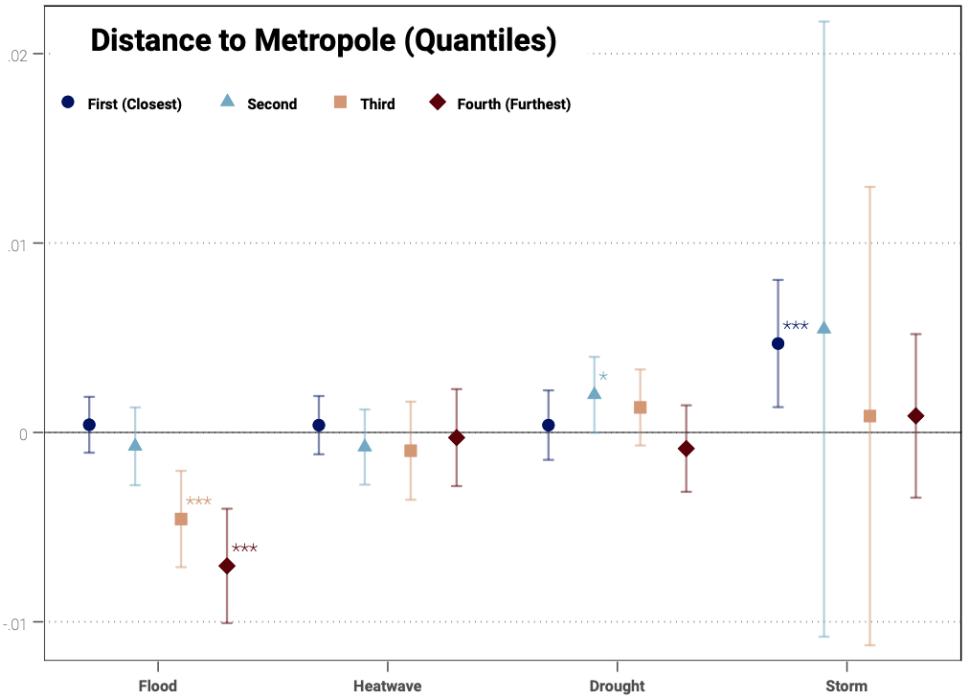

(b) By Distance to Metropole

Fig. 4.4. Heterogeneity Analysis - Part 2/2

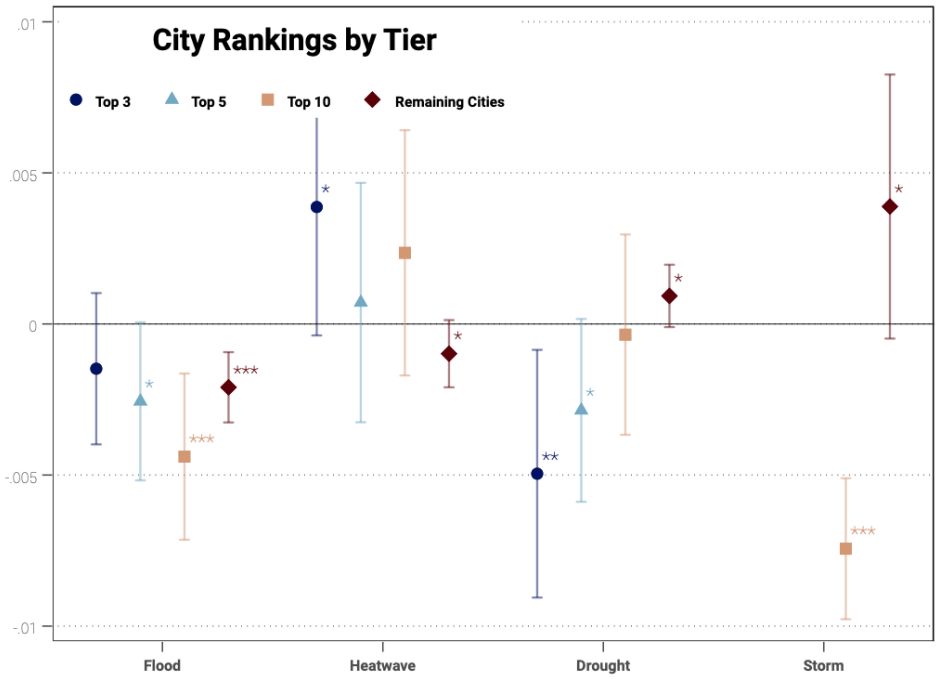

(c) By City Rank

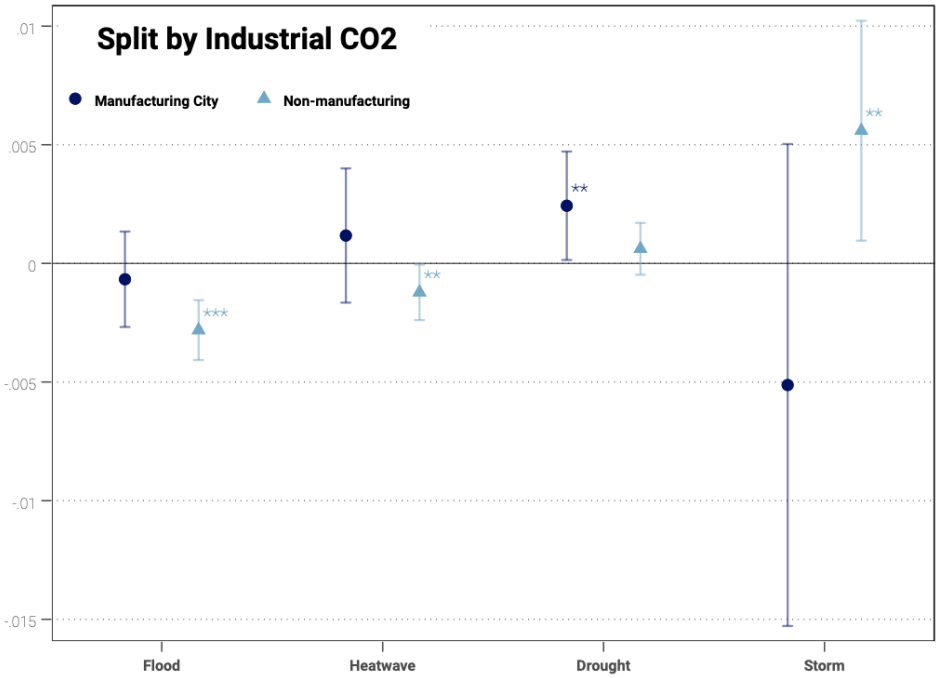

(d) By Industrial Type

Note: The four figures show the results of using Equation 4.1 on stratified samples. Figures show the 95% CIs with p-values *<0.1,**<0.05, ***<0.01. Panel (a) uses stratification by the International Wealth Index (IWI) taken from Lee and Braithwaite (2022). Panel (b) computes the distance of cities to the country's metropole (usually the capital) as defined in OECD and SWAC (2020). Panel (c) ranks the cities by build-up area and creates ranks to determine primary cities in a country. Panel (d) uses carbon emission data from the European Commission - JRC and International Energy Agency (IEA) (2023) to define manufacturing cities as cities with >20% industrial carbon emission over total city emissions.

Table 4.7
Thresholds for Cut-off Intensity (Reinhardt, 2024a)

Hazard	Measure	Onset Type	Share of sample above threshold	Expected Effect on Urbanisation
Heatwave (WBGT)	5-day MA	Slow	2.8%	Mixed
Drought	SPEI-3 & SPEI-6	Slow	1.8%	Positive
Tropical Storm	m/s	Rapid	<1%	Positive
Wind gusts	m/s	Rapid	<1%	Positive
Flood	DFO Index	Rapid	1.3%	Negative

Note: WBGT: Wet Bulb Globe Temperature, MA: Moving Average, SPEI: Standardized Precipitation Evapotranspiration Index, DFO: Dartmouth Flood Observatory. Onset type explains the time it needs for these events to show substantial effects relative to the other included hazards. Original input sources based on: Mistry, 2020; Gebrechorkos et al., 2023; Done et al., 2020; Knapp et al., 2010; Copernicus Climate Change Service, 2024a; Brakenbridge and Dartmouth Flood Observatory, 2023; Prein and Holland, 2018b.

Table 4.8
Descriptive Statistics

	(1) Full sample	(2) Southern Africa	(3) Eastern Africa	(4) Central Africa	(5) Western Africa
Outcomes					
Impervious Area	7.07	14.86	4.89	4.34	6.03
(km^2)	(41.87)	(76.29)	(28.98)	(16.50)	(32.40)
City growth (%)	0.03	0.02	0.03	0.03	0.03
	(0.06)	(0.04)	(0.06)	(0.08)	(0.06)
Treatments					
Flood	0.09	0.10	0.09	0.02	0.12
	(0.33)	(0.32)	(0.32)	(0.15)	(0.38)
Storm	0.00	0.01	0.00	0.00	0.00
	(0.07)	(0.15)	(0.05)	(0.00)	(0.00)
Heatwave	0.11	0.02	0.02	0.06	0.22
	(0.40)	(0.17)	(0.19)	(0.31)	(0.55)
Drought	0.11	0.14	0.12	0.13	0.09
	(0.36)	(0.41)	(0.36)	(0.38)	(0.32)
Other Variables					
High-Risk Flood	14.46	9.48	15.10	16.02	15.68
Areas (% of Urban Area)	(15.81)	(9.73)	(18.05)	(16.04)	(15.93)
Population	903.57	835.28	793.99	795.96	1047.19
Density (per km^2)	(960.33)	(748.30)	(894.08)	(1096.72)	(1011.92)
Absolute Wealth	45.74	60.81	38.42	31.93	49.03
(IWI)	(16.56)	(19.42)	(13.16)	(14.34)	(9.90)
Employment in	52.00	39.70	63.14	64.36	45.50
Agriculture (%) - country	(18.22)	(25.11)	(15.49)	(10.93)	(9.39)
Total GDP per	1575.89	3749.99	672.26	795.61	1485.13
capita - country	(1799.22)	(2803.33)	(396.85)	(1525.07)	(774.10)
Nb. of 10x10km	2.50	2.94	2.91	2.18	2.16
Cells overlap	(3.30)	(3.02)	(4.79)	(2.02)	(2.46)
Nb. Cities[2]	5 721	1 029	1 577	894	2 280
Observations	107796	19551	28260	16716	43254

Note: The means and standard deviations are shown in brackets. [2] shows the number of cities. One might note that the total number of cities does not sum up to the total sample. This is due to some cities being omitted because of a lack of data availability or unrealisticly high growth rates. This table shows the descriptive statistics of an average city in this study. Standard deviations are shown in brackets. Division into Subregions follows the UN Statistical Division. Hazard Exposure indicates the average number of events that occur within any given year.

Table 4.9
Extension: Accounting for Spatial Scope

	(1) 0km	(2) 50km	(3) 100km	(4) 150km	(5) Commuting Zone
∑ Flood Events in t-1	-0.0025*** [-0.003,-0.002]	-0.0030*** [-0.004,-0.002]	-0.0035*** [-0.005,-0.002]	-0.0038*** [-0.005,-0.003]	-0.0046*** [-0.006,-0.003]
∑ Heatwave Events in t-1	-0.0019*** [-0.003,-0.001]	-0.0021*** [-0.003,-0.001]	-0.0019*** [-0.003,-0.001]	-0.0017** [-0.003,-0.001]	-0.0022*** [-0.003,-0.001]
∑ Drought Events in t-1	0.0011** [0.000,0.002]	0.0012** [0.000,0.002]	0.0016** [0.000,0.003]	0.0022** [0.001,0.004]	0.0010 [-0.000,0.002]
∑ Storm Events in t-1	0.0025*** [0.001,0.004]	0.0043*** [0.003,0.006]	0.0051*** [0.003,0.007]	0.0055*** [0.003,0.008]	0.0041*** [0.003,0.006]
Obs.	107,777	107,777	107,777	107,777	99,158
Average Urban Growth	.03	.03	.03	.03	.03
Average Shock by Buffer					
Flood	.066	.066	.066	.066	.063
Heatwave	.213	.215	.217	.217	.222
Drought	.117	.116	.116	.117	.114
Storm	.016	.016	.016	.016	.017
Controls?	Yes	Yes	Yes	Yes	Yes
City-FE	Yes	Yes	Yes	Yes	Yes
Year-FE	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes
Clustered at	City	City	City	City	City

Note: *** 0.01,** 0.05, * 0.1 p-val. This table shows the results of an adjusted version of the baseline model. The description of the shocks and the model are discussed in D.6. Shocks measure here the average fraction of the buffer affected by a shock. Data from commuting zones are based on Data for Good at Meta (2023).

D Appendix to Chapter 4

D.1 Additional Figures

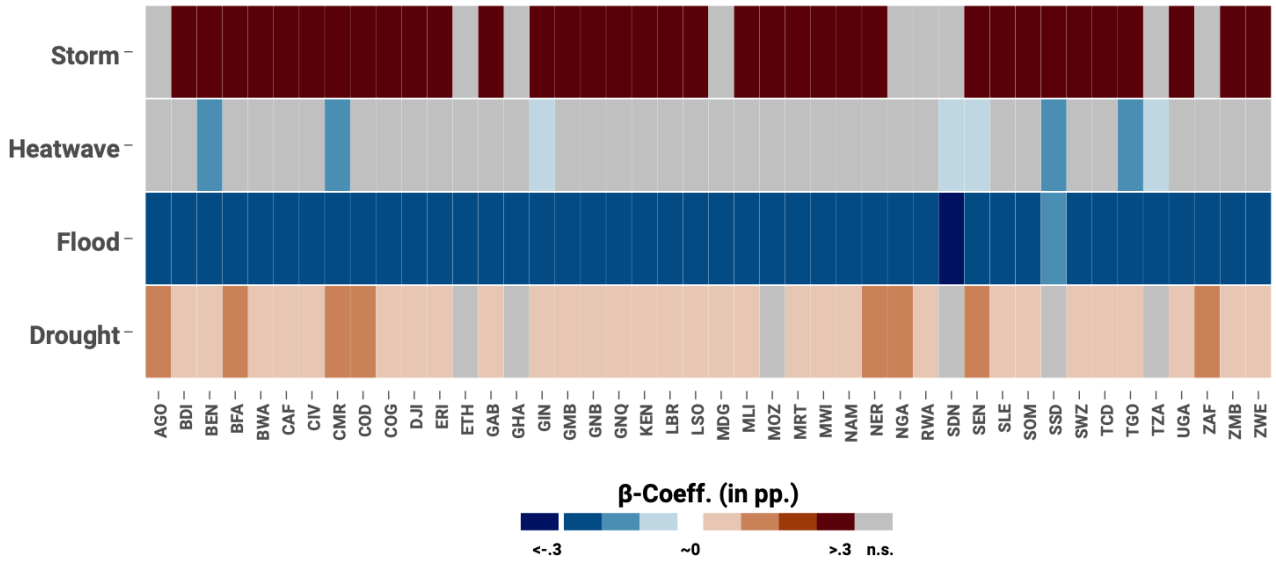

Fig. D.1. Robustness: Leave-one-out-exercise

Note: The figure shows the results of sequentially dropping one country at a time and examining whether the significance level changes for each of the four hazards. The shown confidence intervals are at a 10% level and all regressions follow Equation 4.1 with city, year and region-year fixed effects. Areas in light grey are insignificant at a 10% level.

Fig. D.2. More Information on the Sample

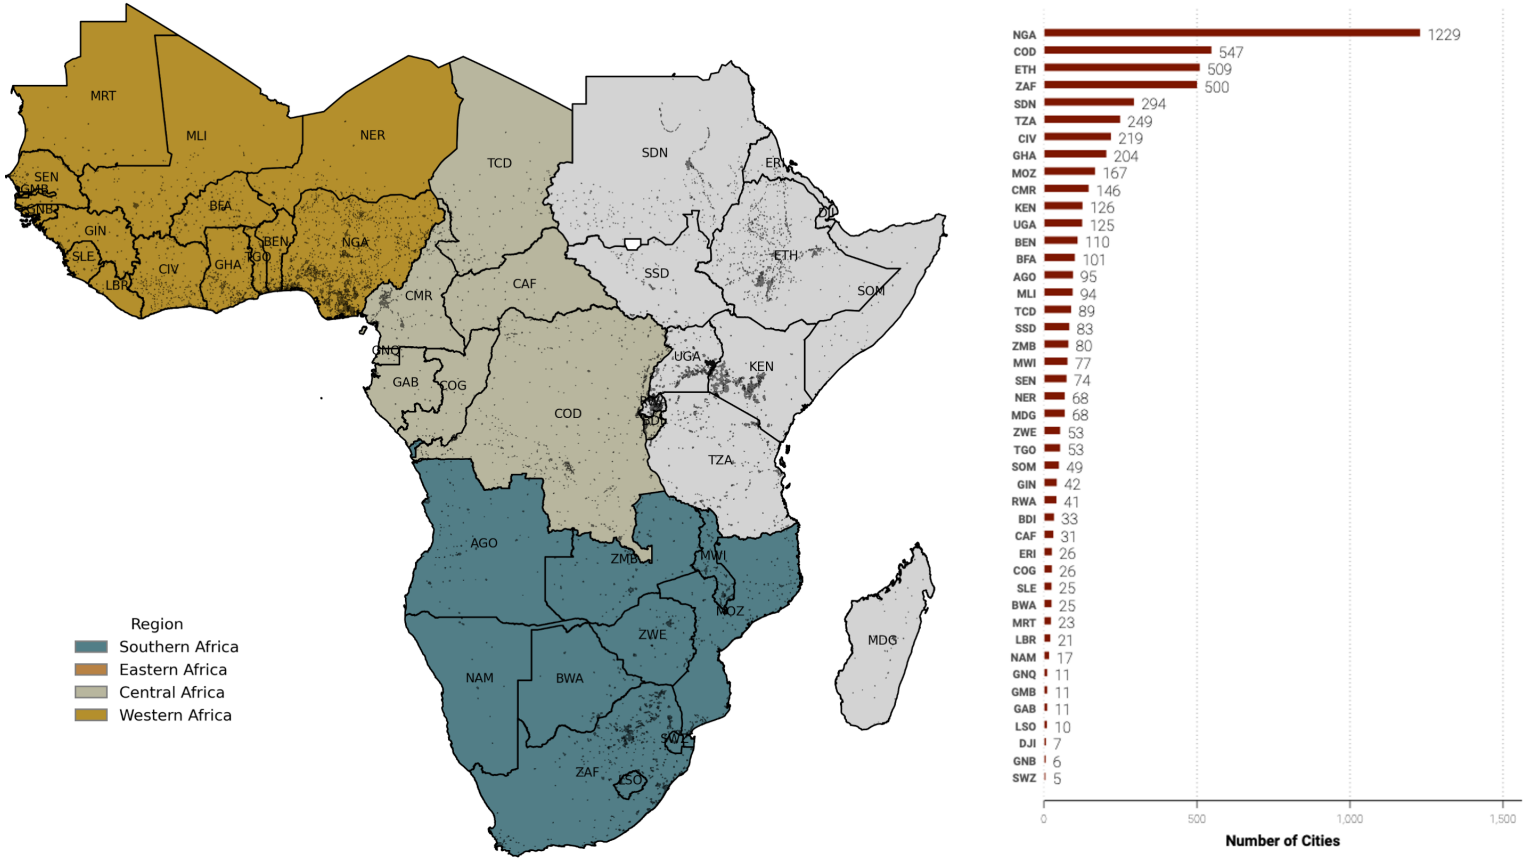

Note: The left figure shows how this paper classifies regions following the UN Statistical Division. The right plots the number of cities by country ISO-3 code. The black dots represent each an urban agglomeration included from OECD and SWAC (2020).

D.2 Additional Tables

Table D.1
Classification of Countries by Region

Region	Countries (ISO3)
Southern Africa	Angola (AGO), Botswana (BWA), Eswatini (SWZ), Lesotho (LSO), Malawi (MWI), Mozambique (MOZ), Namibia (NAM), South Africa (ZAF), Zambia (ZMB), Zimbabwe (ZWE)
Eastern Africa	Djibouti (DJI), Eritrea (ERI), Ethiopia (ETH), Kenya (KEN), Madagascar (MDG), Rwanda (RWA), Somalia (SOM), South Sudan (SSD), Sudan (SDN), Tanzania (TZA), Uganda (UGA)
Central Africa	Burundi (BDI), Cameroon (CMR), Central African Republic (CAF), Chad (TCD), Republic of the Congo (COG), Democratic Republic of the Congo (COD), Equatorial Guinea (GNQ), Gabon (GAB)
Western Africa	Benin (BEN), Burkina Faso (BFA), Côte d'Ivoire (CIV), The Gambia (GMB), Ghana (GHA), Guinea-Bissau (GNB), Guinea (GIN), Liberia (LBR), Mauritania (MRT), Mali (MLI), Niger (NER), Nigeria (NGA), Senegal (SEN), Sierra Leone (SLE), Togo (TGO)

Note: The table shows how the different countries are classified into regions. It follows the classification of the UN (UNSD, 2023)

Table D.2
Robustness: Sample Dependence

	Winsor Urban Growth	Exclude cities that have population (2019) of...			
	(1) Both Tails (5%)	(2) $< 5\,000$	(3) $< 10\,000$	(4) $< 50\,000$	(5) $\geq 50\,000$
$\sum$ Flood Events in t-1	-0.001*** [-0.001,-0.000]	-0.002*** [-0.003,-0.001]	-0.002*** [-0.003,-0.001]	-0.003*** [-0.004,-0.002]	-0.002*** [-0.003,-0.001]
$\sum$ Heatwave Events in t-1	-0.000 [-0.001,0.000]	-0.001 [-0.002,0.000]	-0.001 [-0.002,0.000]	-0.000 [-0.004,0.003]	-0.001** [-0.002,-0.000]
$\sum$ Drought Events in t-1	0.000 [-0.000,0.001]	0.000 [-0.000,0.001]	0.000 [-0.001,0.001]	-0.001 [-0.002,0.000]	0.001** [0.000,0.002]
$\sum$ Storm Events in t-1	0.003*** [0.001,0.005]	0.002* [0.000,0.004]	0.003** [0.001,0.005]	0.012 [-0.002,0.026]	0.003 [-0.001,0.007]
Obs.	107,796	102,190	84,426	21,885	85,911
Average Urban Growth	.02	.03	.03	.03	.03

Note: *** 0.01,** 0.05, * 0.1 p-val. This table shows the baseline specification from Equation 4.1. All models include city, year and region-year fixed effects and standard errors are clustered at the city-level. Population counts are based on data from Liu, Cao, et al. (2024a).

Table D.3
Robustness: Treatment Definition + Non-Linearities

	Events		Time-Decay Events		Number of Months		Intensity	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Flood in t-1	-0.0024***	-0.0025*	-0.0036***	-0.0099***	-0.0008**	-0.0022**	-0.0015***	-0.0010
	[-0.003,-0.001]	[-0.005,-0.000]	[-0.005,-0.002]	[-0.014,-0.006]	[-0.001,-0.000]	[-0.004,-0.001]	[-0.002,-0.001]	[-0.002,0.000]
Heatwave in t-1	-0.0008	-0.0009	-0.0029***	-0.0011	-0.0001	-0.0002	0.0008***	-0.0025*
	[-0.002,0.000]	[-0.003,0.001]	[-0.004,-0.001]	[-0.005,0.003]	[-0.001,0.000]	[-0.001,0.001]	[0.000,0.001]	[-0.005,-0.000]
Drought in t-1	0.0008*	-0.0004	0.0018**	0.0003	0.0004	0.0001	-0.0016***	-0.0007
	[0.000,0.002]	[-0.002,0.001]	[0.000,0.003]	[-0.003,0.004]	[-0.000,0.001]	[-0.001,0.001]	[-0.002,-0.001]	[-0.002,0.001]
Storm in t-1	0.0036*	-0.0040	0.0026	-0.0092	0.0013	-0.0002	0.0006***	0.0015***
	[0.000,0.007]	[-0.012,0.004]	[-0.004,0.009]	[-0.033,0.014]	[-0.002,0.004]	[-0.005,0.004]	[0.000,0.001]	[0.001,0.002]
Flood in t-1 (sq.)		0.0001		0.0075***		0.0005		-0.0002
		[-0.001,0.002]		[0.004,0.011]		[-0.000,0.001]		[-0.001,0.000]
Heatwave in t-1 (sq.)		0.0000		-0.0014		0.0000		0.0001**
		[-0.001,0.001]		[-0.004,0.001]		[-0.000,0.000]		[0.000,0.000]
Drought in t-1 (sq.)		0.0008		0.0015		0.0001		-0.0004
		[-0.000,0.002]		[-0.002,0.005]		[-0.000,0.000]		[-0.001,0.000]
Storm in t-1 (sq.)		0.0046		0.0122		0.0006		-0.0000**
		[-0.001,0.010]		[-0.016,0.040]		[-0.001,0.002]		[-0.000,-0.000]
Obs.	107,796	107,796	107,796	107,796	107,796	107,796	106,732	106,732
Average Urban Growth	.03	.03	.03	.03	.03	.03	.03	.03
Controls?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
City-FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: *** 0.01,** 0.05, * 0.1 p-val. This table shows the baseline specification from Equation 4.1 but with a changed treatment definition. Note that there are fewer observations in Column (1) due to restricting the sample to make it comparable between (1) and (6). Column (1) shows the sample-restricted baseline. Column (3) applies a time-decay weight for the shock. (5) show the number of months in a year in which an event is present. Column (5) uses annual raw intensities. All even columns add quadratic terms. All models include city, year and region-year fixed effects and standard errors are clustered at the city-level.

Table D.4
Robustness: Flood Measure

	Flood Robustness		Has at least one flood of type...			
	(1) Include common flood	(2) Has > 10% flood risk	(3) Any	(4) Pluvial	(5) Riverine	(6) Both
Has Flood	-0.0018*** [-0.003,-0.001]	-0.0017** [-0.003,-0.000]	-0.0024*** [-0.003,-0.001]	-0.0137*** [-0.018,-0.010]	-0.0027*** [-0.004,-0.002]	-0.0026 [-0.006,0.001]
Obs.	107,796	107,796	107,796	100,749	103,757	107,796

Note: *** 0.01,** 0.05, * 0.1 p-val. This table shows the baseline specification from Equation 4.1. Column (1) changes the treatment definition from floods being above DFO Index 2 to DFO Index 1 so account for weaker floods. Column (2) adopts a more restrictive definition, where floods only count if at least 10% of an urban area is at high flooding risk. Column (3) uses the definition of the baseline with the change that it is binary (has a flood or not), independent of the number. (4) defines a pluvial flood if the month in which flood occurs has a rainfall anomaly of 2 standard deviations based on data from Climate Hazards Center (UCSB) (2023). (5) defines a flood if in same month a 3-day streamflow anomaly is detected with at least 2 standard deviations above historical means based on Joint Research Center and Copernicus Emergency Management Service (2019). (6) is defined if both conditions hold. Missings are higher in (4)-(6) as years where other flood type occurred are ommitted. All models include city, year and region-year fixed effects, controls and also heatwave, drought and storm events and standard errors are clustered at the city-level.

D.3 Mathematical Construction of the Shock

This section explains mathematically how the 10x10 km grid of Reinhardt (2024a) is adjusted on a monthly level to align with city boundaries, enabling the computation of a yearly city panel dataset.

At first, the 10x10km grid overlaps with the city polygons and overlapping areas are computed for each cell. Then, a city-specific weighted mean for each hazard is computed based on the following equation:

$$\overline{x}_{imt}^{(k)} = \frac{\sum_g A_{ig} \cdot x_{gmt}^{(k)}}{\sum_g A_{ig}}$$

where $\overline{x}_{imt}^{(k)}$ is the weighted average of hazard intensity for hazard k in city i, month m and year t. It is computed as summing, across all all 10x10km grid cells g, the product of two components: (1) the area overlap between city i and cell g denoted as A_{ig} ("*What part of a grid cell overlaps with city boundaries?*") and (2) the cell specific intensity value for the hazard ($x_{gmt}^{(k)}$). This is then divided by the total overlap area of all grids within a city. This intensity can be the 5-day maximum wet-bulb temperature (heatwave), the SPEI-03 and SPEI-12 (drought) or the maximum windspeed recorded (storm). This is divided by the total number of overlapping area. For floods, the maximum recorded DFO Index is used.

Next, the monthly intensities are converted into binary dummies by applying the predefined thresholds.

$$haz_{imt}^{(k)} = \begin{cases} 1 & \text{if } \overline{x}_{imt}^{(k)} > \tau^{(k)} \\ 0 & \text{otherwise} \end{cases}$$

where $\tau^{(k)}$ are the defined thresholds and the weighted average needs to be above the threshold. Every hazard has a specific defined τ described in section 4.2. Then, the annual city-level event number is defined as the number of events that switch from 0 to 1.

$$haz_{it}^{(k)} = \sum_{m \in M \setminus \{\min(M)\}} \mathbb{1}(haz_{it,m-1}^{(k)} = 0 \wedge haz_{it,m}^{(k)} = 1)$$

where m is the month of the year as part of M, the total number of months, but excluding the first month of the panel (January 2000) as the previous month could not be determined. This will give us the total number of "unique" events for hazard k in city i in year t. The final vector Haz contains four elements ($haz_{it}^{(k)}$), one for each hazard.

Especially for flood, this approach might overlook events. If for instance flood 1 lasted from February to March and another flood occured in March to April, this would be considered as a single event. To avoid this, for flood, the total number of events is compared to the unique counts of flood-specific identifiers of floods that intersect with a city.

D.4 Additional Heterogeneity Tests

By Geographic Region and ENSO period This additional analysis tests if the coefficients presented in Table 4.3 vary across timing in the El Niño-Southern Oscillation (ENSO). It is included that, depending on the cycle, within the ENSO, extremely wet or extremely dry conditions can be expected. The effect of ENSO varies by broader regions, which is why it fits well in the geographic heterogeneity analysis. We suspect that during ENSO-neutral years, extreme events (like floods or droughts) may have stronger or different effects. This is because they are less anticipated, and thus may catch governments or communities more by surprise than when expectation and preparedness are higher. To define ENSO we take the Oceanic Niño Index (ONI) from the NOAA Climate Prediction Center (2025) that allows to define months as part of neutral, El Niño or La Niña months. El Niño or La Niña years are defined if at least four consecutive months of a year fulfill conditions according to the ONI to be classified as El Niño or La Niña month. The dummies are added as interaction terms with the shocks in a stratified region of the following form:

$$\begin{aligned}\Delta \log(Y_{icrt}) = \alpha_r + \sum_{k=1}^{4} \beta_{k,r} \, \text{Haz}^{(k)}_{icr,t-1} + \sum_{k=1}^{4} \delta_{k,r} \, \text{Haz}^{(k)}_{icr,t-1} \cdot \text{Niño}_{t-1} \\ + \sum_{k=1}^{4} \theta_{k,r} \, \text{Haz}^{(k)}_{icr,t-1} \cdot \text{Niña}_{t-1} + \gamma_r X_{c,t-1} + \lambda_{t,r} + \lambda_{i,r} + \varepsilon_{icrt}\end{aligned}$$

The results are shown in Table D.5. In the sample analyzed, urban expansion in Eastern Africa does not appear to be significantly influenced by climate hazards during El Niño or La Niña years. Instead, observed effects are primarily associated with hazards occurring during ENSO-neutral years. For Southern Africa, floods have a significant impact on urban growth, though the direction varies by ENSO phase: floods during neutral years reduce urban expansion, whereas during La Niña years, they are estimated to increase urban growth. Droughts in this region reduce urban expansion, but this effect is significant only during non-neutral ENSO phases. In Central Africa,

droughts during El Niño or La Niña years are associated with reduced urban growth, effectively offsetting any positive effects observed during neutral years. For Western Africa, the negative association of floods and urban expansion of Table 4.3 is likely driven by floods occurring during neutral ENSO years. Interestingly, during ENSO years, floods may be associated with increased urban growth, suggesting a potential context-specific or adaptive response. Heatwaves also show negative effects on urban expansion but only in ENSO years. Lastly, droughts also have negative effects in neutral years, but they also show positive effects. A possible explanation for the stronger responses in general during neutral years is the lower expectation of extreme weather events during such periods. As a result, early warnings might show larger uncertainties. Unlike El Niño or La Niña phases, which are often accompanied by specific forecasts and preemptive responses, neutral years may result in less preparedness, thereby amplifying the strong responses of particularly droughts and floods.

D.5 Constructing of the Compound Shocks

In this part of the paper the previously individually regarded shock metrics are disentangled into two components: number of events of a hazard type k_j where no other hazard than k_j occurred in temporal and spatial proximity and those where another hazard type $k_{l \neq j}$ occurred. Furthermore, two increasing constraining restrictions are included: first, just account for a temporal window in which k_l needs to take place and second, the order of sequence in which k_j and k_l take place.They will be presented as *unordered* and *ordered* sequences, respectively.

Mathematical Intuition Behind the Unordered Sequence

First, we construct a compound event in an unordered sequence as happening in the same month or within a given time frame. To this end we construct monthly moving windows that define if within that period both type of hazards are present or not. Having computed on the city-level the presence of a hazard, we obtain the following (note e is an event of a hazard type and element of the earlier presented vector haz):

$$e_{im}^{(fl)}, e_{im}^{(hw)}, e_{im}^{(dr)}, e_{im}^{(st)} \in \{0, 1\}$$

So for all of the four hazards (flood, heatwave, drought and storm) we know if the month exceeded a threshold or not. Note here as we currently study the effects over the years the subscript t is omitted as the windows, which will be introduced next, move across years. Next for a rolling-window of a specified length and for each hazard, we

Table D.5
Heterogeneity by Region and ENSO Cycle

	(1) Southern Africa	(2) Eastern Africa	(3) Central Africa	(4) Western Africa
In neutral years				
∑ Flood Events in t-1	-0.0046** [-0.008,-0.001]	0.0000 [-0.003,0.003]	-0.0010 [-0.013,0.011]	-0.0078*** [-0.010,-0.006]
∑ Heatwave Events in t-1	-0.0043** [-0.007,-0.001]	0.0040** [0.001,0.007]	-0.0020 [-0.006,0.002]	0.0003 [-0.001,0.002]
∑ Drought Events in t-1	0.0017 [-0.000,0.004]	0.0036*** [0.002,0.006]	0.0081* [0.000,0.016]	-0.0052*** [-0.007,-0.004]
∑ Storm Events in t-1	0.0046 [-0.001,0.010]	0.0033 [-0.000,0.007]	0.0000 [0.000,0.000]	0.0000 [0.000,0.000]
In El-Niño Years..				
∑ Flood Events in t-1	0.0045 [-0.001,0.010]	0.0028 [-0.001,0.007]	-0.0097 [-0.023,0.003]	0.0041*** [0.002,0.007]
∑ Heatwave Events in t-1	0.0202 [-0.019,0.060]	-0.0010 [-0.006,0.004]	-0.0033 [-0.008,0.001]	-0.0020** [-0.004,-0.001]
∑ Drought Events in t-1	-0.0052*** [-0.008,-0.002]	0.0025 [-0.000,0.005]	-0.0094* [-0.018,-0.001]	0.0038*** [0.002,0.006]
∑ Storm Events in t-1	0.0014 [-0.002,0.005]	-0.0031 [-0.008,0.001]	0.0000 [0.000,0.000]	0.0000 [0.000,0.000]
In La-Niña Years...				
∑ Flood Events in t-1	0.0069*** [0.003,0.011]	0.0007 [-0.003,0.005]	0.0049 [-0.011,0.020]	0.0049*** [0.002,0.008]
∑ Heatwave Events in t-1	0.0062 [-0.006,0.018]	0.0064 [-0.023,0.036]	-0.0012 [-0.023,0.020]	-0.0049*** [-0.008,-0.002]
∑ Drought Events in t-1	-0.0023* [-0.005,-0.000]	-0.0004 [-0.004,0.003]	-0.0120** [-0.020,-0.004]	0.0096*** [0.005,0.014]
∑ Storm Events in t-1	-0.0036 [-0.011,0.004]	-0.0020 [-0.010,0.006]	0.0000 [0.000,0.000]	0.0000 [0.000,0.000]
Obs.	19,551	28,275	16,716	43,254
Average Urban Growth	.02	.03	.03	.03
Controls?	Yes	Yes	Yes	Yes
City-FE	Yes	Yes	Yes	Yes
Year-FE	Yes	Yes	Yes	Yes
Region-Year FE	No	No	No	No

Note: *** 0.01,** 0.05, * 0.1 p-val. Confidence intervals are at a 90% level. This table shows the baseline results of Equation 4.1. All samples include controls (a country's employment in agriculture (in %) and the countries total GDP per capita) and have city and year fixed effects. Regions are defined as subregions of Africa and are shown in Figure D.2. El Niño and La Niña are defined if at least 4 month of a year fulfill these conditions taken from NOAA Climate Prediction Center (2025).

define the rolling maximum:

$$\tilde{e}_{i,m}^{(j)}(w) = \max_{h \in W_w(m)} e_{i,h}^{(j)}, \quad W_w(m) = \{m - \lfloor (w-1)/2 \rfloor, \ldots, m + \lceil (w-1)/2 \rceil\}$$

Which gives us $\tilde{e}_{i,m}^{(j)}(w)$ as the rolling-maximum of hazard k with window length w in city i and month m. It is defined as maximum value over the month-block h which is part of a new set of aggregated month-sets W. This set W just defines the months-windows to consider when computing the shock. In our sample we define $w \in \{1, 2\}$, so 1- or 2-month windows. Having computed the block-maxima for each of the four hazard types, we build a raw compound hazard indicator for each unordered pair (j,l), whereby $j \neq l$ between both types:

$$cp_{i,m}^{(j,\ell)}(w) = 1\left(e_{i,m}^{(j)} = 1 \wedge \tilde{e}_{i,m}^{(\ell)}(w) = 1 \quad \vee \quad e_{i,m}^{(\ell)} = 1 \wedge \tilde{e}_{i,m}^{(j)}(w) = 1\right)$$

This will give us $cp_{i,m}^{(j,\ell)}(w)$ as one of six possible combinations of j and l. The criteria requests that either j is 1 and falls into a rolling maxima of l or the other way round. So for the example of flood and heatwaves we essentially ask:

$$cp_{i,m}^{(j,\ell)}(w) = 1 \text{ iff } \begin{cases} \text{flood in } m \text{ AND heat-wave somewhere in window w,} \\ \text{OR} \\ \text{heat-wave in } m \text{ AND flood somewhere in window w.} \end{cases}$$

To define compound events, we want to avoid double-counting consecutive months. For example, if a heatwave and flood combination extend over two months, we want to define it as a single event rather than two separate events. Thus, we define a final monthly equation that only keeps the first month of the spell:

$$haz_{i,m}^{(j,\ell)}(w) = 1\left(cp_{i,m}^{(j,\ell)}(w) = 1 \wedge cp_{i,m-1}^{(j,\ell)}(w) = 0\right)$$

So here we only keep the first occurrence of a specific sequence. In a final step for each hazard combination ($haz_{i,m}^{(j,\ell)}(w)$) with a set window w we compute the number of events per year. To avoid double accounting for the same type of shock, we subtract the number of compound events from the total number of each event involved. For instance, there were two flood events recorded and one heatwave-flood pair. The included number for flood will be one, given that presumably the other flood overlapped with the heatwave event. This model is not perfect as there might be some risk that using the "take-the-first-month" spell condition, might overlook other newly occuring

events. This means that we likely underestimate the number of events and the measure is likely a lower bound estimate.

Empirical Estimation of the Unordered Sequence

To only account for the presence of a second hazard type independent of the order between hazard type k_j and k_l the baseline equation Equation 4.1 is thereby adjusted to accomodate this division:

$$\Delta \log(Y_{icrt}) = \alpha + \sum_{k=1}^{4} \beta_k \cdot \widetilde{\text{Haz}}^{(k)}_{icrt-1} + \sum_{c=5}^{11} \theta_c \cdot \text{CP}^{(c,w)}_{icrt-1} + \gamma X_{ct-1} + \lambda_t + \lambda_i + \lambda_{r\times t} + \epsilon_{it} \tag{4.2}$$

where $\widetilde{\text{Haz}}^{(k)}_{icrt-1}$ is similar to Equation 4.1 but excludes from the count of hazard events, events where within a window w another hazard type has been recorded, thus the tilde. It is therefore on average lower than our baseline variable. New is the part of $\text{CP}^{(c,w)}_{icrt-1}$ that is a variable that counts the number of compound combinations of type c in a given window w. c is one of six potential unordered pairs (e.g., flood-heatwave), and the window is set in the model to either one or two months. The interpretation of θ is the marginal short-term effect of having one additional compound shock of type c in the last year on the growth rate of urban expansion in percentage points.

Mathematical Intuition Behind the Ordered Sequence

As before here we have defined the four distinct hazard types per month-event:

$$e^{(fl)}_{im}, e^{(hw)}_{im}, e^{(dr)}_{im}, e^{(st)}_{im} \in \{0, 1\}$$

Every month our algorithm investigates if none, one, or more than one hazard occurred. Formally, we build a hazard set S for each month-city pair that accounts for the combinations:

$$S_{im} = \left\{ k \in \{\text{FL}, \text{HW}, \text{ST}, \text{DR}\} : e^{(k)}_{im} = 1 \right\}$$

Which is for instance for a month with flood and heatwave ("FL;HW") in a month under consideration (m). If nothing is reported then $S_{im} = \emptyset$. Next, we define a forward-looking window (w), in our case set to two months. Essentially, if the considered month is m, we look forward into the months m+1 and m+2.

For the same city i consider any pair of months (m,n) with $m < n \leq m + w$. Thus, m is the current months and n is a later month inside the forward looking window w. For example, if the current month is March (m=3), the window extends to May (m+w=5),

and n would be April (4). Then, we build ordered pairs as

$$k_1 \rightarrow k_2, \quad k_1 \in S_{im}, k_2 \in S_{in}, k_1 \neq k_2$$

Note here that k_1 is defined at month m and k_2 is defined at month n. As a next step, to account for unique sequences in a given year, we introduce a city-specific memory set that checks if the hazard sequence within a given window is already recorded or not. We name this memory set $\mathcal{P}_i$ and check if either $k_1 \rightarrow k_2$ or $k_2 \rightarrow k_1$ already exists within the set. If the sequence or its reverse is already recorded, we ignore it. Otherwise, we add it to $\mathcal{P}_i$ at the later month n (as we presume this is where the sequence is fully established and shows the most significant effects). This condition filters out repeated sequences or order reversals that might take place due to fading out events (an area sees a reduction of flooded areas while experiencing still drought conditions).
We enforce a 12-month cool down period. This means if a compound sequence is recorded at a particular months, the same sequence (in either order) cannot be recorded again within the next 12 month. This reduces concerns of double-counting the same compound combination. Formally, a sequence is then defined as:

$$\text{Seq}_{in} = \begin{cases} & \exists m \text{ s.t. } n - w \leq m < n, \\ k_1 \rightarrow k_2 & k_1 \in S_{im},\ k_2 \in S_{in},\ k_1 \neq k_2, \\ & (k_1 \rightarrow k_2) \notin \mathcal{P}_i(n-1),\ (k_2 \rightarrow k_1) \notin \mathcal{P}_i(n-1) \end{cases} \tag{4.3}$$

Essentially Seq_{in} records a hazard-sequence $k_1 \rightarrow k_2$ if neither $k_1 \rightarrow k_2$ nor $k_2 \rightarrow k_1$ has occured in the previous 12-month window. If these conditions are satisfied, the new sequence is added to $\mathcal{P}_i$ (at month n). We look at $\mathcal{P}_i(n-1)$ because the condition set must use a memory set before month n is updated, otherwise if we use $\mathcal{P}_i$ at time n, the new sequence would already be added to $\mathcal{P}_i$ and thus disqualify itself.
Once the sequence variable (Seq_{in}) is defined, we obtain unique counts by city i and hazard-sequence $k_1 \rightarrow k_2$[25]. As the variable is currently on a monthly level, we aggregate each unique sequence into a city-year panel. In short, for each city we identify hazard k_1 (e.g. flood) occuring in a given month, then look forward within a window (w) to check if another hazard k_2 (e.g. heatwave) occurs. If so, we form an ordered pair $k_1 \rightarrow k_2$. If this pair (or its reverse $k_2 \rightarrow k_1$) occured in the past 12 months, we exclude

[25] For simplicity reasons, here only k_1 to k_2 pairs are shown. In practice, the designed algorithm in Python also allows for more than just one ordered pair to enter the sequence. In practice thus we capture more than just one bivariate pair.

it; otherwise, it is labelled as a new compound sequence[26].

Empirical Estimation of the Ordered Sequence

The derived empirical estimation is fairly similar to the previous one:

$$\Delta \log(Y_{icrt}) = \alpha + \sum_{k=1}^{4} \beta_k \cdot \widetilde{\text{Haz}}^{(k)}_{icrt-1} + \sum_{j=5}^{J} \theta_j \cdot \text{CPseq}^{(j)}_{icrt-1} + \gamma X_{ct-1} + \lambda_t + \lambda_i + \lambda_{r\times t} + \epsilon_{it} \tag{4.4}$$

Here $\text{CPseq}^{(j)}_{icrt-1}$ is a vector of the j-th ordered pair of the vector CPseq defined for a fixed 2-month window. All else remains the same.

How does the number of combinations change between models?

The change in definition between a month, a two-month unordered setup, and an ordered sequence with a 12-month cooling period will change the number of existing combinations. The hazard-pair counts are shown in Figure D.3. One can see that for all hazards, the number of events increases, allowing for a larger time window. The anti-diagonal is empty as we exclude the possibility of within-hazard compound events[27]. The most common combination of hazards is heatwaves and floods, and heatwaves and droughts, followed by floods and droughts. Storms, given the spatial concentration in some parts of Eastern and Southern Africa, see substantially fewer compound events. It is not surprising to see that heatwave-flood pairs are the most common type given that they might be connected through convective storms[28]. For an ordered 2-month period, we see that the by far most common sequence in Sub-Saharan Africa is heatwaves followed by flood events. The next most common are droughts followed by floods and droughts followed by heatwaves. Note that the number of events is not identical to the unordered 2-month period due to slightly different computations of the sequence mentioned above.

[26] Tested but not included is using other thresholds of the "exclusion criteria" of 12 months. Allowing for shorter windows increases the noise and risk of double accounting, but includes more event combinations. Broadly comparable results have been documented for a 10 month and 14 month window.

[27] Within-month within-hazard compound events would require higher temporal resolution, which we do not have. Otherwise, within the year, the adjusted $\text{Haz}^{(k)}_{icrt-1}$ accounts for the number of events of that type in a given year

[28] A possible example for this hazard pair could be in the Sahel zone. In May and June, before the northward shift of the Intertropical Convergence Zone (ITCZ) and the onset of the West African Monsoon (WAM), the region experiences intense heat, potentially sufficient to be considered heatwaves. When the monsoon arrives in July, the combination of moist air from the Gulf of Guinea and the superheated land surface leads to strong atmospheric instability. This can trigger intense convective storms, sometimes resulting in floods.

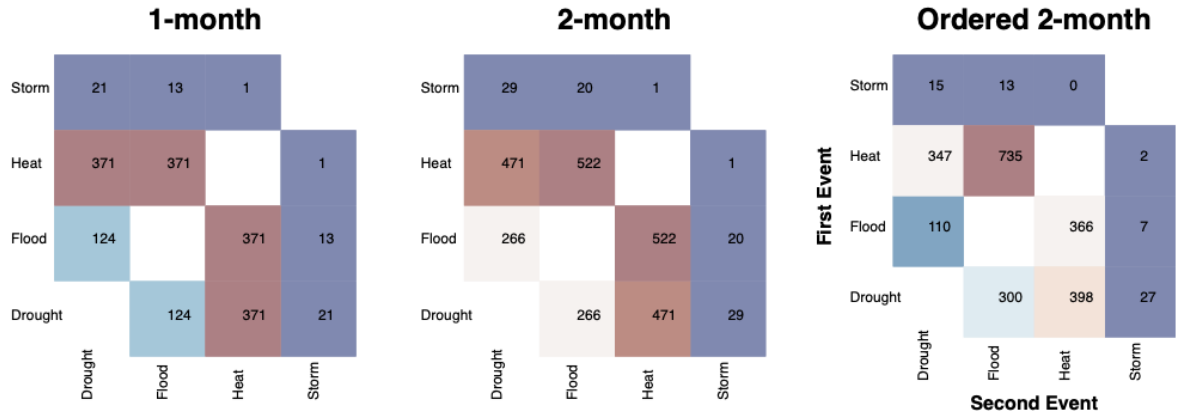

Fig. D.3. Compound Sequence: Number of Hazard Pairs

Note: The figure shows the number of unordered (1-month and 2-month) and ordered hazard pairs. The creation of these pairs follows the methods described above. Note that for the ordered pairs, the y-axis indicates the first, the x-axis indicates the second event.

D.6 Constructing Spatial Shock

In an extension of the baseline model of Equation 4.1 we include the exposure to flood, storm, drought and heatwaves also in nearby rural land. To this end, we employ two approaches to assess weather shocks in nearby land: (1) a spatial buffer analysis around each urban agglomeration and (2) an analysis based on the commuting zones of each major city. The following sections describe the methodology used to compute weather shocks under each approach.

Creation of the Buffers

First, the spatial buffer approach draws a 50km, 100km, and 150km catchment area around the centroid of each urban agglomeration. The different buffers allow one to assess the sensitivity of the results to different magnitudes of buffer. An illustration of these buffers can be seen in Figure D.4 (a) and (a1) for the example of Kinshasha. As the treatment variables of drought, heatwave, storm, and flood are taken from Reinhardt (2024a), the same data is used to compute the number of shocks within these buffers. The 10x10km grid of Reinhardt (2024a) is overlapped with the buffers and intersections are computed. If buffers exceed national boundaries, they are clipped to only include areas within the same country. In Figure Figure D.4 (2) one can see that for Kinshasha only the catchment area within the Democratic Republic of Congo is included, while Congo-Brazzaville in the Republic of the Congo is excluded. This is done under the assumption that if internal migration happens, migrating into a bor-

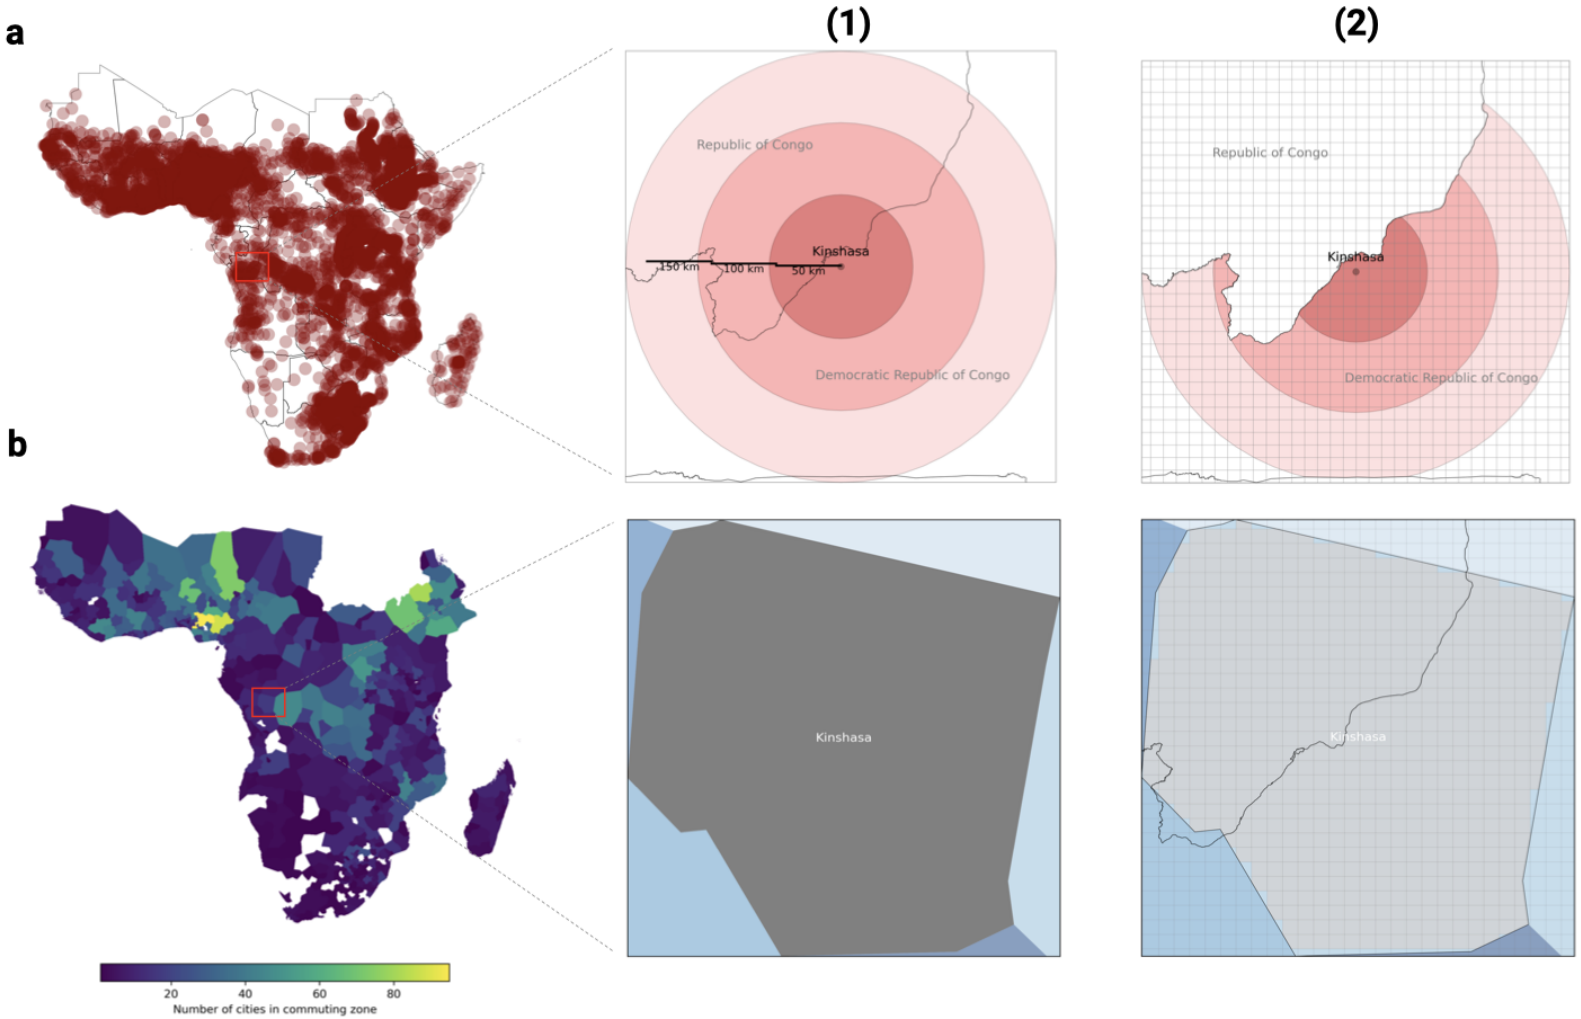

Fig. D.4. Illustration of Creation of the Spatial Shocks

Note: The figure shows in Panel (a) the buffer-approach to compute shocks. In (1) a 50,100 and 150km buffer is drawn around each city (here in Kinshasha, Democratic Republic of Congo). In (2) the buffers are overlayed with the 10x10km grids based on Reinhardt (2024a) and city-cells are excluded and in (3) additionally cells with less than 10 people are excluded. In panel (b) the commuting-zone approach is shown. First, it shows the number of cities and the commuting zones included in the sample. In (1) the commuting-zone for "Kinshasha" is plotted, in (2) the grid as mentioned above is added and finally in (3) urban areas and uninhabitated cells are excluded. The remaining cells that are not a city and have more than 10 people living there are used to compute the number of shocks.

der country is costly, and in this case excluded by construction. Having defined the buffers, a *continuous* shock variable is calculated for each buffer. The computation is slightly different from that of the baseline model. We modify the treatment definition to more closely reflect the spatial distribution of the shocks. Instead of relying solely on the urban agglomeration, we construct spatial buffers (e.g., 50km) around urban cores and intersect these with the 10x10km grid. Here, each grid cell is directly converted to a binary shock indicator using the same threshold as before. These are then summed annually[29]. The new shock is the average within each buffer[30]. For cities within these buffers, the same method is applied to all grid cells intersecting the urban shapefile, ensuring comparability. As a result, the new shock variable becomes continuous (average area of a buffer affected), reflecting the intensity of shocks within the buffer area. An increase in this measure can be interpreted as, for example, a 100% increase in the number of shocked grid cells within the buffer. This differs from the baseline, where the focus is on the number of events derived from average intensities within an urban area. The baseline does not focus on the average number of shocks in cells with which the city (or its buffers) intersects. Given the description above, a problem arises if the number of cells used to compute the buffer might vary (e.g. Kinshasa has only half of the cells given that the other half falls into another country). To account for this, additionally, the total population within each buffer is computed using data from GlobPop (Liu, Cao, et al., 2024a).

Static buffers have two main limitations. First, they do not account for country-specific contexts such as the size of a country. For instance, a 150 km radius covers a significant

[29] Following the notation of D.3, we compute

$$hexp_{gmt}^{(k)} = \begin{cases} 1 & \text{if } x_{gmt}^{(k)} > \tau^{(k)} \\ 0 & \text{otherwise} \end{cases}$$

so computing for each month and cell g if thresholds are exceeded. Then sum them as

$$hexp_{gt}^{(k)} = \sum_{m \in M \setminus \{\min(M)\}} \mathbb{1}\left(hexp_{gt,m-1}^{(k)} = 0 \wedge hexp_{gt,m}^{(k)} = 1\right)$$

This gives the total number of events for type k in cell g

[30] In the notation of D.3 this would correspond to

$$hexp_{ict}^{(k,b)} = \frac{1}{G_i^{(b)}} \sum_{g \in G_i^{(b)}} hexp_{gt}^{(k)}$$

, where G is the total number of grid cells in buffer b to city i. The value of G might vary between cities due to the overlap of the boundaries.

portion of a small country like Burundi, whereas it covers only a minor part of a much larger country like Niger. Second, buffers ignore terrain and infrastructure, which are important to assess given that they are strongly related to the costs of displacement and mobility. A 150 km journey through the hilly and densely forested Kivu region in the Northeastern Democratic Republic of Congo is more time-consuming and potentially expensive than an equivalent distance in the well-connected catchment area of Cape Town, South Africa. To account for this, an alternative buffer scheme is used: commuting zones provided by Data for Good at Meta (2023). The commute zones are based on daily movement maps from home locations to potential work locations of individuals who have activated the location services of used Meta apps (such as Facebook). By tracking weekday movements, the project attempts to build commute zones to nearby population centers (*nodes*) and isolate them by drawing Voronai shapes around these catchment areas. These data were taken for all African countries considered in this article and intersected with urban agglomerations. The 416 included commuting zones that cover 5 185 cities of the 5 821 cities in the full sample are shown in Figure D.4 panel (b).

The commuting zones are named after the main destinations of weekday movements. I identify the main commuting node in each commuting zone using fuzzy matching on the name and based on a distance-matching criterion. This is important, as many, especially smaller urban agglomerations, fall within the commuting zones of larger cities. For example, the small agglomeration of Bankana, 120km south-east of Kinshasa, in the Democratic Republic of Congo, falls into the Kinshasa commuting zone. This means that our final sample will have in parts multiple cities in the same commuting zone with one city being the "population hub". Similarly to the buffer approach explained above, the 10x10km grids as in the main approach are overlapped with each commuting zone and 10km grid cells that contain urban agglomerations and population less than 10 people are excluded (see lower row (2) and (3)). Attentive readers may notice that the commuting zone of Kinshasa in Figure D.4 extends across the river into the Republic of Congo, including parts of Brazzaville. This is intentional, as the analysis is based on functional urban catchment zones rather than strict administrative boundaries. If shocks drive migration to larger urban areas, some internal migrants in the Republic of the Congo may relocate to Brazzaville instead of Kinshasa, both of which are captured in the sample. Similarly to the buffer approach, based on the same threshold and methodology as in the main paper, flood, drought, storm, and heatwave events are computed for each cell. Finally, the resulting yearly events are then aggregated on

a city level as the maximum number of events and the average number of events in the catchment area, e.g the average number of flood events recorded in 2015 in the Kinshasa catchment areas (excluding areas with high urban presence and absence of local population) for all cities that fall into this catchment area.

Computation of Model

To estimate the effects of a 100% shock in the catchment area (buffer or commuting zone) a slightly modified version of the baseline model is used:

$$\Delta \log(Y_{icrt}) = \alpha + \sum_{k=1}^{4} \beta_k \cdot \text{Hexp}_{icrt-1}^{(k,b)} + \gamma X_{ct-1} + \kappa \log(\text{Pop}_{icrt-1}^{(b)}) + \delta_t + \lambda_i + \lambda_{r \times t} + \epsilon_{it} \tag{4.5}$$

The model is comparable to the original specification, with three key exceptions. First, instead of the vector Haz, which contained categorical indicators for each of the four types of hazards, we now use the vector Hexp, which measures continuous hazard exposure for each hazard. Second, this exposure vector is defined at a specific buffer size b, chosen from one of the following: 0 km, 50 km, 100 km, 150 km, or the commuting zone. The interpretation of each coefficient β also changes: rather than capturing the effect of an additional discrete shock, each β now reflects the effect of an additional shock affecting the full buffer. Lastly, we introduce a new control variable: the lagged log of population within buffer b. This accounts for the size of the potentially affected population and improves comparability across buffers that may contain a different number of grid cells due to exclusions near borders. Including lagged population along with hazard exposure risks that effects of weather shocks on urbanisation are mediated through population. This risks to underestimate the effects if population endogenously respond to the hazards, which likely in the case of weather shocks. Including it nonetheless here serves three purposes: first, it reduces risk of ommitted variable bias. Second, it improves comparability of effects across catchments with varying baseline population levels. Third and lastly, in includes indirectly a key mechanism that may drive urban expansion: rural-urban migration. For these reasons, we add population directly in the same lag structure into the model.

D.7 Constructing the Event-Study

Overview

This section provides more context on the use of the event-study estimator of panel matching suggested by Imai and Kim (2019). This estimator is especially useful in cases

with non-random treatment assignment and repeated treatments over time. Following Baker et al. (2025), there are currently only very few causal difference-in-difference-like estimators that allow treatment to switch from treated to nontreated, an assumption that is crucial in the case of extreme events. In addition, the estimator proposed by Imai and Kim (2019) is relatively flexible, as it allows conditioning on both time-varying and time-invariant covariates, unlike most other estimators, with the exception of the one by Liu, Wang, and Xu (2024). Furthermore, standard two-way fixed effects models rely on strong identification assumptions, notably strict exogeneity: no time-varying confounders (i.e., no omitted variable bias), no carry-over effects from past treatments, and no feedback from past outcomes to current treatment. The method of Imai and Kim (2019) instead relies on the assumption of sequential ignorability, which is less restrictive. This assumption allows for carry-over effects from past treatments and outcomes onto the current outcome, but assumes that the current treatment is not influenced by past covariates. Although sequential ignorability is generally easier to defend—especially in the context of heatwaves—there remains a small risk of violation. This risk arises if, for example, urban expansion increases the likelihood of heatwaves. However, we argue that while heatwave risk and heatwave occurrence may be correlated, the latter depends on additional stochastic factors. These include, among others, precipitation, local weather variability, and local adaptation efforts. Therefore, the increase in risk does not directly translate to an increase in heatwave events. Based on this reasoning, we consider the sequential ignorability assumption plausible in our setting and proceed using the estimator by Imai and Kim (2019).

The Setup

The method proposed follows a three-step procedure. First, a matched set is created for each transition of the shocks. Note here that floods, heatwaves, droughts and storms are considered binary, given that the panel match estimator does not allow for categorical treatments. Thus, for every city, a set of matched cities with similar treatment histories for an event is iteratively created. For example, if Nairobi was flooded in 2008 and 2015, the algorithm searches for other cities (in our case, we set the size of the matching group to 5) that were flooded in 2008 but not in 2015. Second, once the set of cities with similar treatment histories has been selected, a more comparable control city is created by refining the matched set using matching methods. In our case, we tested for Mahalanobis matching, covariate balance propensity score matching and propensity score weighting. To build a better fitting sample, decisions need to be made on which covariates the matched sample should be refined. Given that

the outcome is urban expansion, the covariates include the last two years of all other hazards except the one in question. This assumes that their past treatment might influence urban expansion. Furthermore, we use the initial urban area in 1990 to get a better approximation on the growth level the cities have seen in the past. Finally, we take the last four years before a treatment in terms of urban expansion to get cities with similar historical growth rates. However, it is not trivial to decide which matching procedure to use.We assess the best-fitting model by comparing the average difference of variables used for matching between the treated unit and the weighted average of the control units over all matched sets. This assessment follows the author's guidelines. The results are shown in Figure D.5. The general standard deviation was the smallest for covariate balancing, making it our preferred specification for this study. Having

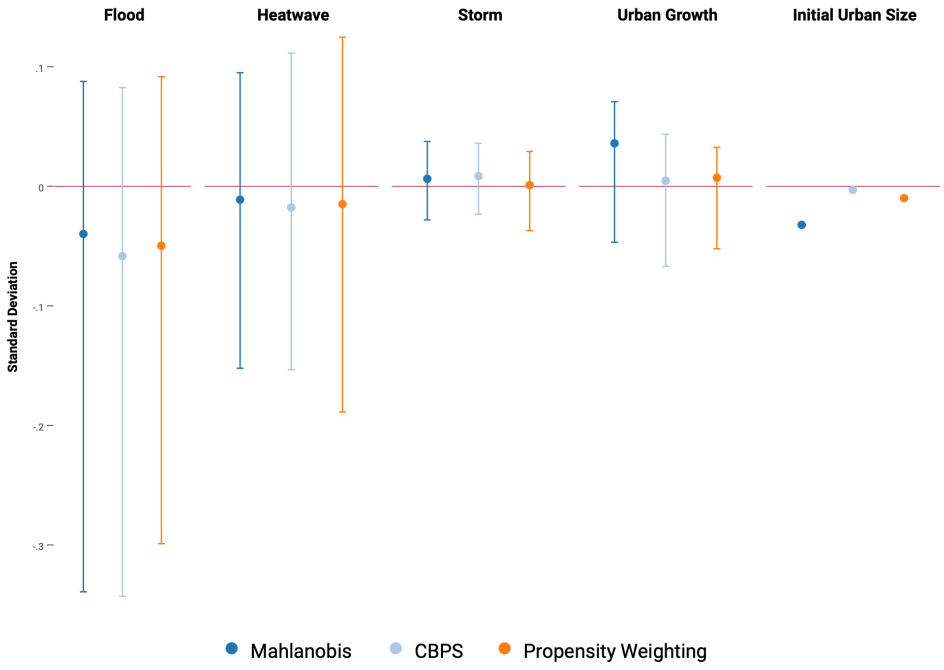

Fig. D.5. Event-Study: Selecting Best Matching Method

Note: The plots compare the difference in covariate balancing for the three tested matching methods. The closer coefficients are to the zero line the better on average the matching performance. Outcome is the standard deviation for the average difference between treated units and average matched units.

selected the pretreatment window for the treatment history, the matching method and covariates to match on, in the third step, the average treatment effect of treated (ATT) is computed between the treated unit and the refined set. Standard errors are computed using block bootstrapping with 100 replications.

Additional Results

The results of all different methods are shown in Figure D.6. For flood, all models show a similar pattern. Heatwaves show only significant positive effects when using propensity score weighting. A reason might be that propensity score weighting does not exclude poor matches, but might be more model-dependent than Mahalanobis or Covariate Balance Propensity Score matching. For drought and storm, again, similar results are found. We also test, yet not shown, using different pre-treatment lags and different lag structures in the covariates.

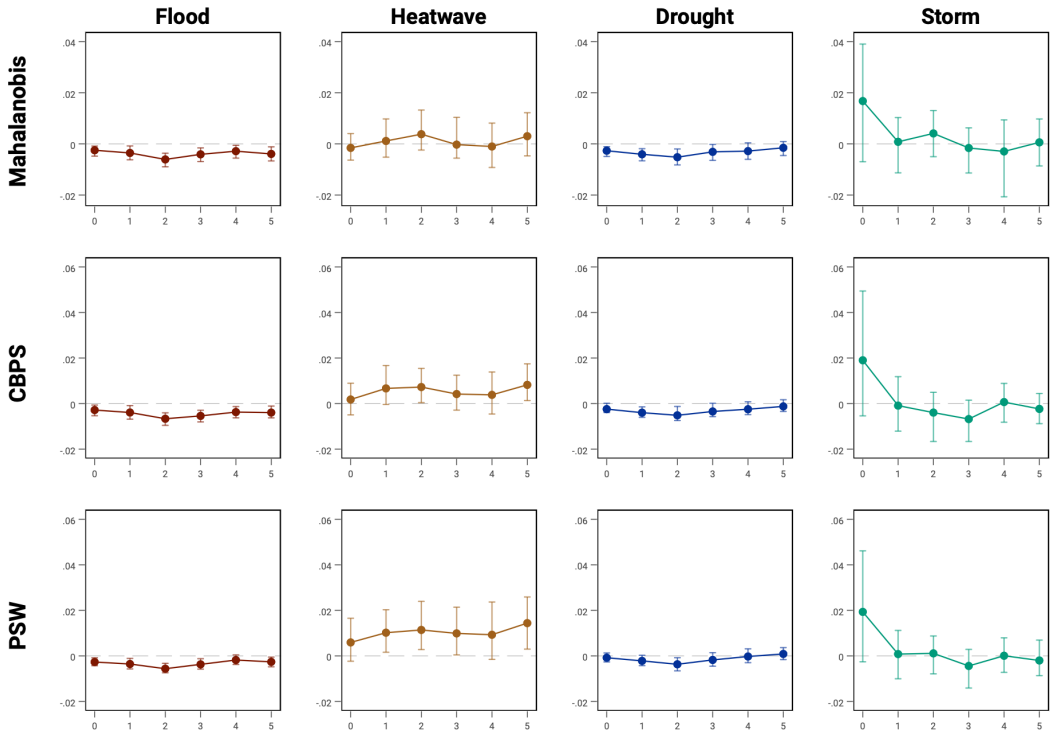

Fig. D.6. Event-Study: All Models

Note: This figure shows the results of using panel matching by Imai and Kim (2019) on the four different hazard on physical urban expansion. Treatment histories are fitted for an eight-year window and standard errors are bootstrapped 100 times. The confidence intervals shown represent a 10% level. The first row shows Mahalanobis matching, the second Covariate Balance Propensity Score matching and the last row shows propensity score weighting.

Robustness

To test whether the results are driven by strong pre-trends, a placebo test is conducted to assess whether parallel trends hold. The results are shown in Figure D.7 and appear to hold for all but heatwaves, where one period is significant. Thus, the mild and

mostly insignificant effects of heatwaves should be interpreted with caution.

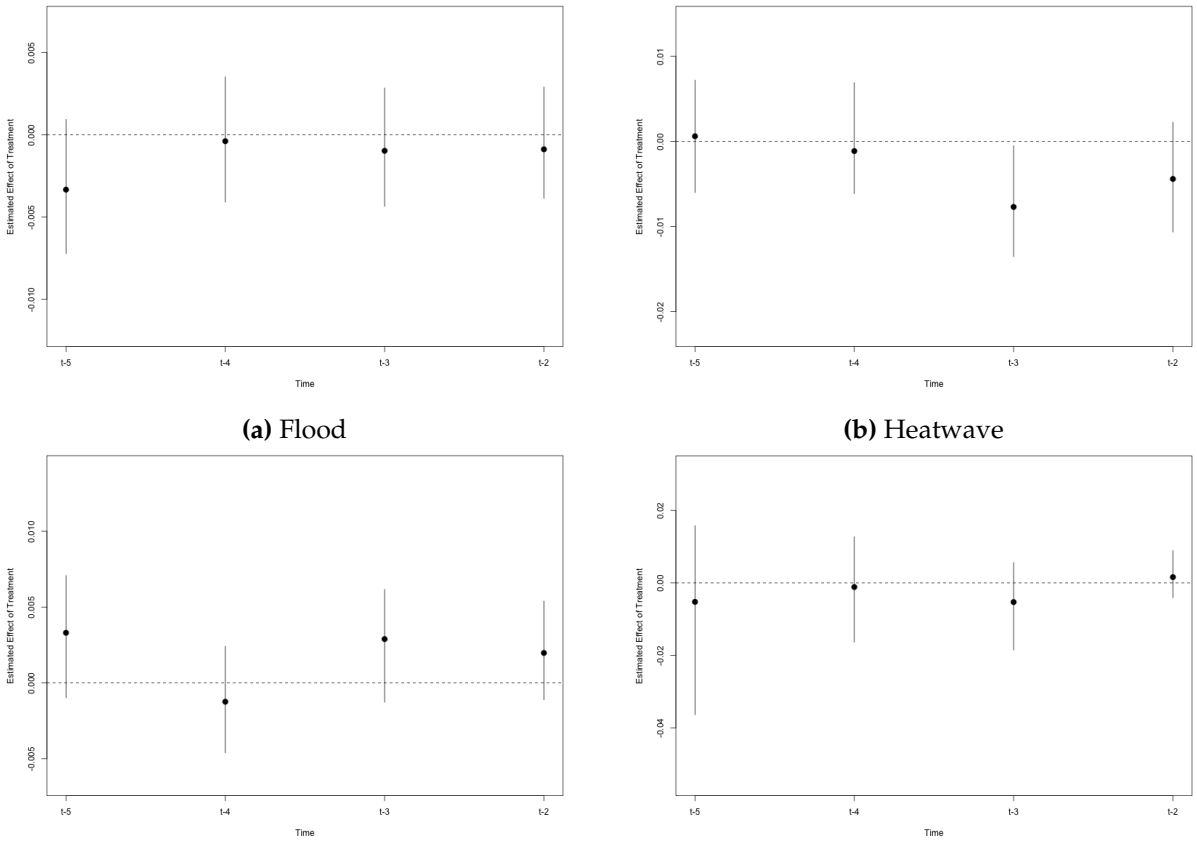

(a) Flood **(b)** Heatwave

(c) Drought **(d)** Storm

Fig. D.7. Event Study: Placebo Tests

Note: The four panels show the placebo test on the pre-treatment sample using the same covariate setup as in the main event-study. Estimates are based on the tests proposed by Imai and Kim (2019). Confidence intervals shown represent the 10% level. The model tested here is the Covariate Balance Propensity Score Model.

D.8 Within-City Analysis: Data and Transformation

This section describes the data used and transformations applied to estimate the results of Table 4.6.

Data Sources

Population As mentioned in the main part of the paper, the data are based on a dataset by Liu, Cao, et al. (2024a). This data provides annual population counts and density estimates at a resolution of 1km. Most annual data for population counts is interpolated from census data and fine-tuned with remote sensing extensions (e.g. building foot-

prints, nightlight, etc.). The data by Liu, Cao, et al. (2024a) seems to perform equally well or slightly better than comparable annual population datasets.

Informal Settlements Acquiring data on informal settlements is challenging, since the criteria that define them differ. They can generally be defined by high land tenure insecurity, lack of access to roads and urban infrastructure, little to no amenities (such as hospitals, canalization system) or areas with high social marginalization of vulnerable groups (UNHCR, 2024). Here, we rely on the Million Neighborhoods project (Bettencourt and Marchio, 2023) to focus on the lack of access to street networks as a measure of "informality". The data used to define these areas is recent, thus the assumption is made that areas with lack of access today will have likely had little to no access in the part. The data provided come in a range zero up to 67 in our sample, where neighborhoods with values above 10 are classified as "off-network" and above six have a very low connectivity. Thus, we use 10 as a strong "informal" metric and 6 as a "weak" informal metric. Within each of these areas, the population is computed using a downscaled estimate for population from Liu, Cao, et al. (2024a), described below. Instead of relying on urban expansion by impervious surfaces here, we focus on population given that impervious surfaces might underestimate the extent of informal activity (knowing that they also capture increase in infrastructure beyond just buildings.)

Flood Risk While already briefly presented in the Data Section, flood risk data is drawn from Rentschler (2024) and coastal flooding risk from Vousdoukas et al. (2018) and Vousdoukas (2023), both have an ultra-high resolution of less than 300m. High-flood risk zones are defined as areas that would be inundated by more than 50 cm in the event of a 100-year flood.

Conflict There exist two common databases for conflict: the ACLED databases (Raleigh, Kishi, and Linke, 2023) and the Uppsala Conflict Data Program (UCDP) (Sundberg and Melander, 2013). We choose the former and account for conflict only if it took place in the form of violent conflicts or riots within a city's urban boundary. However, the dependent outcome variable is also tested for using a binary dummy (whether it has conflict or not) and using the UCDP dataset. The results are comparable.

Downscaling of Population Data

Attentive readers might have noticed that the population data presented comes at a 1km resolution while flood risk as well as neighborhood data might be in a substantially smaller geographic area. To compute population estimates, we apply a simple

downscaling method[31]. In order to downscale the data, we use building footprints on a 100-m resolution from the Global Human Settlement Layer (GHSL) dataset (Pesaresi, 2023). To account only for population locations, only the residential part of building footprints is used. The motivation is that the population will be allocated where the residential build-up is recorded. Although this approach is rather simplistic, GHSL has been used in other literature also as one of several factors besides interpolated population count and Landsat images (see e.g. Tziokas et al., 2023). As the GHSL data is only available in 5-year periods, we always use the last available year. After bilinear resampling and coarsening, downscaling weights are computed based on the normalized weight of the 100m area to the total sum of build-up within the 1km grid. To reduce noise from individual data points, a Gaussian smoothing filter is applied. We apply here for the following model

$$G(x,y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2+y^2}{2\sigma^2}\right)$$

, where G() is the Gaussian function at each lon-lat pair (x,y), a σ value of 1.The Gaussian filter spreads cell-specific estimates over local neighbors, reducing the possibility of extreme outliers. The sigma of one is a rather small amount of smoothing as we want to retain a sufficiently high variation in population. Applying the Gaussian smoothing might lead to a slight underestimation of population serving rather as a lower bound estimate.

To compute area-specific statistics, we rely on an auxiliary 300m grid within each city that defines if an area is informal or not or if it is at high flood risk or not (see panel (c) and (d) of D.8 for a visualization). Once downscaling is completed, zonal statistics are then computed with the sum of population within this auxiliary grid. Given the centroid overlapping approach implemented in the *rasterstats* package in Python, the sum values might be slightly overestimated for population, which is why the means are also computed for robustness. An illustration of the procedure is shown in D.8. To benchmark the estimates, population estimates are validated for three points in time in three countries with unconstrained interpolated population data of 100 m adjusted with UN estimates of WorldPop (2020). The results are shown in Figure D.8. In comparison with the World Pop data, the downscaled data slightly overestimates the population

[31] We are aware that annual population estimates exist at higher resolution, but using them for all of Africa exceeded the scope of this paper. However, we validate our estimates against them on a random selection of cities.

of cells. The R^2 varies between 0.46 and 0.69. Although that indicates not a perfect prediction in comparison with other interpolated data sets, given that large fractions of the overestimation is due to outliers that are reduced as in the main specification, the outcome is going to be logged, overall, the matching quality is sufficiently good.

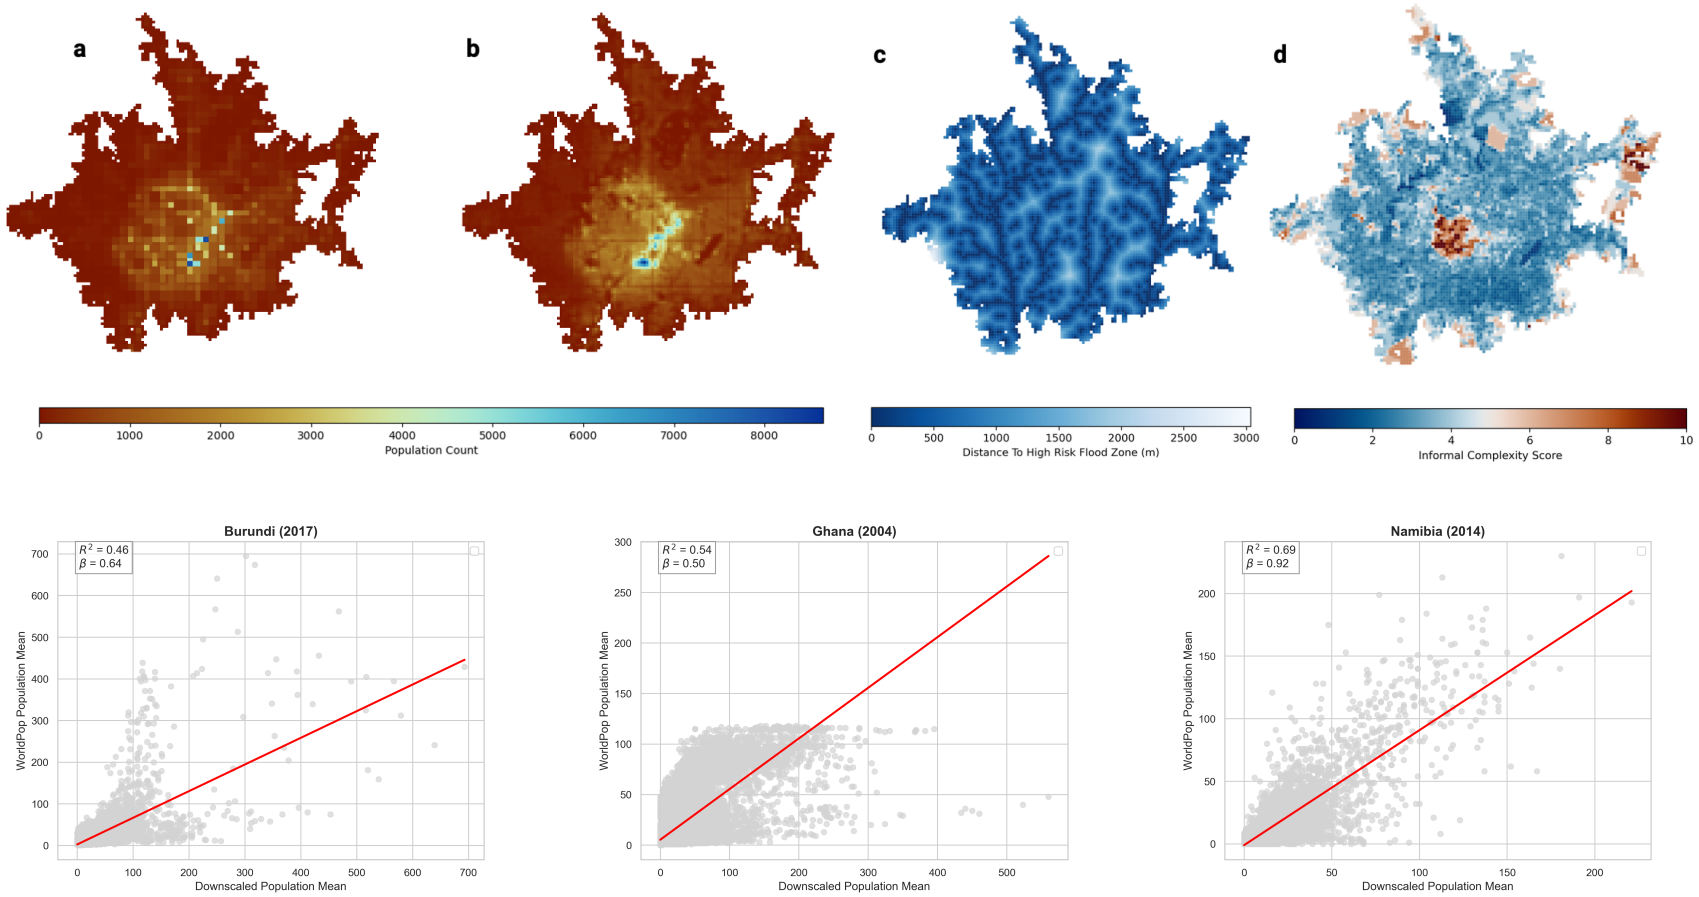

Fig. D.8. Downscaling of Population data and Robustness Tests

Note: The figure visualizes the data generating process. The upper panel shows the data. Panel (a) presents the original population count data at a 1 km resolution. (b) provides the downscaled version using the produre described in the text. (c) visualizes high flood risk zones (dark blue where distance to flood zone is <300m) and (d) shows the used informality complexity measure by are. The lower panel shows for three countries and years the validation of a linear regression of the 100m World Pop data from WorldPop (2020) with the 100m downscaled version based on GHSL residential building footprint. Color palettes follow Crameri, Shephard, and Heron (2020).

Résumé de la thèse

Le changement climatique anthropique influence non seulement le climat dans lequel nous vivons, mais aussi, et surtout, les interactions que les êtres humains entretiennent avec la nature qui les entoure. Le changement climatique, classé parmi les neuf limites planétaires par Rockström et al. (2009), représente un seuil critique à l'intérieur duquel l'humanité peut évoluer en toute sécurité. En tant que l'un des défis les plus importants et les plus largement débattus du XXIe siècle, il souligne la nécessité d'une coopération mondiale pour maintenir la stabilité du système terrestre. Ne pas y parvenir constitue non seulement une menace pour la prospérité économique, mais compromet également les objectifs de développement déjà atteints, comme ceux formulés dans les Objectifs de Développement Durable (ODD) des Nations Unies. Depuis la découverte des gaz à effet de serre par les travaux de Fourier, Tyndall et Arrhenius au XIXe siècle, un nombre croissant d'études et d'articles ont vu le jour à partir du milieu du XXe siècle, avec l'établissement de la science climatique moderne (pour un aperçu, voir Dobes, Jotzo, and Stern, 2014). Depuis lors, la communauté internationale a mis en place divers accords internationaux, notamment le Protocole de Kyoto et, plus tard, l'Accord de Paris. Des projets de collaboration ont également été établis, à l'instar du Groupe d'experts intergouvernemental sur l'évolution du climat (GIEC) ou du Fonds vert pour le climat. Ces initiatives ont pour objectif de limiter l'ampleur du phénomène et d'informer sur ses effets climatiques néfastes. Dans les médias et la littérature scientifique, le discours sur le changement climatique est souvent présenté de manière dichotomique : en se concentrant soit sur l'atténuation, qui vise à limiter l'ampleur et la vitesse du réchauffement climatique, soit sur l'adaptation, qui traite des ajustements face à ses conséquences. Cette thèse se concentrera sur ce dernier aspect, en mettant particulièrement l'accent sur ses implications pour les pays en développement.

L'adaptation au changement climatique est un terme générique qui englobe toute une série de stratégies. Celles-ci comprennent entre autres le développement d'infrastructures résilientes, de systèmes agricoles adaptés et de services de santé capables de faire face aux chocs climatiques. Les économistes ont joué un rôle essentiel dans le débat scientifique entourant ces stratégies, notamment en ce qui concerne l'analyse coût-bénéfice des politiques ou l'évaluation des risques économiques liés au climat. Il est d'une im-

portance capitale de comprendre le risque auquel une personne, une entreprise, une région ou un pays est exposé face au changement climatique lors de la conception et de l'évaluation de stratégies d'adaptation.Selon le Groupe d'experts intergouvernemental sur l'évolution du climat (GIEC), les risques liés aux extrêmes peuvent être analysés à travers trois dimensions fondamentales. La première concerne l'*exposition*, c'est-à-dire la mesure dans laquelle une population ou un territoire est situé dans une zone touchée par un événement. La deuxième est le danger, qui renvoie à la probabilité qu'un tel événement se produise. Enfin, la *vulnérabilité* désigne la sensibilité d'une région ou d'un groupe à cet événement, une fois qu'il survient (cf. IPCC, 2012, p.69). Parfois, on y ajoute aussi la notion de *réponse* (« comment les individus ou les zones réagissent-ils à l'événement ? »). Les articles issus de ma thèse, présentés ci-dessous, cherchent à combiner différents aspects du risque climatique. Ils articulent notamment la notion d'*exposition*, en analysant la fréquence des événements extrêmes, et celle de *vulnérabilité*, en s'intéressant à certaines des régions les plus pauvres du monde. D'après le Groupe d'experts intergouvernemental sur l'évolution du climat (GIEC), un large consensus existe au sein de la communauté scientifique quant à l'augmentation attendue des événements climatiques extrêmes — tels que la puissance des tempêtes tropicales ou la variabilité des précipitations — à l'échelle mondiale. Toutefois, cette intensification variera selon les régions géographiques. Les pays à revenu faible et intermédiaire sont davantage affectés par les événements climatiques extrêmes, en raison de contraintes financières sévères et d'une résilience plus faible liée à des niveaux de pauvreté élevés, à une gouvernance défaillante et à des conflits persistants. Cette vulnérabilité est illustrée par la Figure 1.

Figure 1 montre la distribution mondiale de l'exposition aux événements extrêmes en fonction du niveau de développement humain. Les données sont désagrégées à une résolution spatiale de 10×10 km, sur la période allant de 2004 à 2015[1]. Les données sont issues du Chapitre 2 de cette thèse de doctorat combinées avec celles sur l'Indice de Développement Humain (IDH). En moyenne, les zones en rouge sont des zones avec un indice de développement humain relativement plus faible par rapport au reste du monde. Il apparaît clairement que de nombreuses régions du Sud, no-

[1] Pour la définition des événements extrêmes, les seuils suivants sont pris : vague de chaleur (WBGT >28), inondation (indice DFO >1), incendie de forêt (>20 % brûlés), tempête tropicale (>33 m/s), tempête de vent (>30 m/s), grêle (>1 jour avec Prob >90%), tremblement de terre (>5 MMI), invasion acridienne (nombre >1) (selon les données de Food and Agriculture Organization (FAO) of the United Nations, 2023), tempête de poussière (>1 jour), glissement de terrain (nombre >1), tsunami (vague >0m), volcan (VEI >0), pollution de l'air (>75 µg/m^3)

tamment l'Afrique, l'Asie du Sud et certaines parties de l'Amérique du Sud, présentent des niveaux relativement faibles d'Indice de Développement Humain, mais des niveaux plus élevés d'exposition observée aux événements extrêmes. Les ménages les plus pauvres étant déjà les plus vulnérables au changement climatique (voir par exemple Hallegatte, Vogt-Schilb, Rozenberg, et al., 2020), il apparaît dès lors primordial de mettre l'accent sur les pays à revenu faible et intermédiaire.

Fig. 1. Distribution mondiale de l'exposition aux événements extrêmes selon l'Indice de Développement Humain (2004–2015)

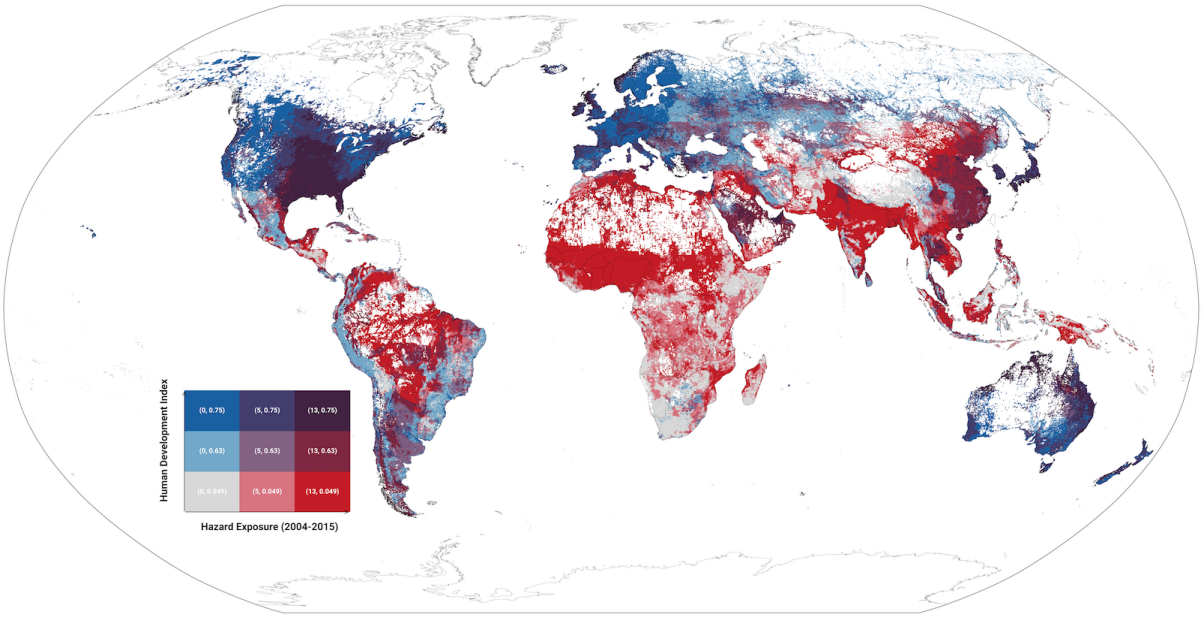

Note : Cette figure est basée sur l'ensemble de données compilé dans Reinhardt (2024a), combinant les 19 risques climatiques les plus fréquents et les données de l'Indice de Développement Humain de Sherman et al. (2023). La figure montre la distribution mondiale de l'exposition aux événements extrêmes en fonction de l'Indice de Développement Humain (IDH) relatif. Les seuils utilisés sont spécifiés en blanc. Seules les zones comptant plus de 10 personnes par kilomètre carré sont incluses.

Pourtant, une question essentielle se pose : comment peut-on relier les événements extrêmes aux résultats économiques ? Afin de connecter les conditions climatiques aux résultats associés à l'économie et à d'autres disciplines, le Groupe d'experts intergouvernemental sur l'évolution du climat (GIEC) introduit le concept de CID – les facteurs d'impact sur le climat – soit des composantes climatiques spécifiques qui affectent directement les systèmes humains ainsi que les écosystèmes (Seneviratne, Nicholls, et al., 2012). Ces CID sont généralement évalués à l'échelle locale ou régionale, là où les

interactions entre les humains, les écosystèmes et les conditions climatiques sont les plus marquées. L'influence des conditions climatiques et des événements extrêmes sur les résultats socio-économiques constitue un fil conducteur de cette thèse. Bien que ces phénomènes ne soient pas toujours désignés explicitement sous le terme de climate impact drivers, plusieurs chapitres s'appuient sur eux de manière fonctionnelle, en particulier en les intégrant comme variables de traitement dans les modèles empiriques. Parmi les exemples notables, on peut citer l'utilisation de la température et des précipitations (« pluviométrie ») comme variables explicatives pour analyser leurs effets sur des dimensions telles que les systèmes agricoles (Aragón, Oteiza, and Rud, 2021; Deschênes and Greenstone, 2007; Lobell, Schlenker, and Costa-Roberts, 2011), les villes et les établissements humains, ou encore le thème de migration (Beine and Jeusette, 2021; Bohra-Mishra, Oppenheimer, and Hsiang, 2014; Cattaneo, Beine, et al., 2019) et la santé (Meierrieks, 2021; Watts et al., 2017). De manière plus générales, les principales variables concernent aussi la croissance économique et la production (Nordhaus, 2019; Tol, 2018; Dell, Jones, and Olken, 2012; Burke, Hsiang, and Miguel, 2015). Beaucoup de ces études ont été citées des centaines voire des milliers de fois. D'autres travaux se sont concentrés plus spécifiquement sur les effets d'un type particulier de choc climatique sur la dimension microéconomique des ménages ou des entreprises. Parmi les exemples récents notables, on trouve l'étude de l'impact des incendies de forêt sur la mortalité (Burke, Driscoll, et al., 2021) ou encore l'effet des cyclones tropicaux sur les actifs des entreprises (Pelli et al., 2023).

Un thème récurrent dans ces études est qu'elles se concentrent généralement sur un type spécifique de choc climatique, tel que la température, les cyclones tropicaux ou les anomalies de précipitations, afin de mesurer l'impact sur une dimension de résultat particulière. Seules quelques études empiriques prennent explicitement en compte les interactions entre différents types de climate impact drivers ou, plus généralement, d'événements extrêmes (voir, par exemple, Fezzi and Bateman, 2015). Dans les modèles théoriques tels que les Integrated Assessment Models, les interactions sont intégrées, mais ces modèles restent souvent insuffisants pour capturer toute la complexité des interactions réelles.

Simplifier l'impact des événements extrêmes sur les résultats, notamment dans les articles d'économie appliquée, en les réduisant à des événements isolés permet aux chercheurs de diminuer la complexité des effets liés aux dangers et de mieux établir des liens de causalité. Cependant, les apports des climatologues, géoscientifiques ou météorologues montrent que cette simplification néglige fréquemment la nature com-

plexe du système climatique. La littérature scientifique sur les événements extrêmes en climatologie a évolué vers une vision plus holistique, tenant compte des interactions complexes entre les phénomènes extrêmes individuels. En particulier, la notion de connected hazards ou de compounding hazards[2] a récemment suscité un vif intérêt. Les paragraphes suivants reviendront sur ce concept et préciseront son importance dans le cadre de cette thèse.

Tout au long de cette dissertation, chaque chapitre défend l'idée que les résultats socio-économiques ne sont pas uniquement influencés par des risques individuels. Les dangers isolés ont tendance à interagir avec d'autres événements extrêmes ou avec des vulnérabilités préexistantes. Cette thèse couvre un éventail plus large de dangers que ceux généralement discutés, bien que la liste des dangers considérés reste non exhaustive.

Événements connectés, leur importance pour la littérature économique et défis méthodologiques

Avant d'aborder les défis méthodologiques posés par les dangers connectés, une brève introduction stylisée est présentée sur ce que sont les événements connectés ou composés, ainsi que sur les raisons pour lesquelles ils devraient interpeller les économistes.

Selon le Groupe d'experts intergouvernemental sur l'évolution du climat SREX, les événements composés désignent des phénomènes climatiques ou météorologiques qui se manifestent simultanément, peu de temps les uns après les autres, ou dans une proximité spatiale, et qui tendent à entraîner des effets amplifiés par rapport aux événements individuels (cf. Seneviratne, Nicholls, et al., 2012, p.118). En complément, en utilisant la typologie proposée par Zscheischler, Martius, et al. (2020), les événements composés peuvent être classés comme suit : les événements pré-conditionnés (une condition climatique amplifie l'impact d'une autre), les événements multivariés (deux ou plusieurs climate impact drivers ou dangers se produisent simultanément), les événements composés dans le temps (deux événements ou plus surviennent dans une séquence rapprochée), ou les événements composés dans l'espace (deux événements ou plus se produisent dans une proximité géographique). Chaque combinaison spécifique entre type d'événement et type de composition peut entraîner des ef-

[2] Dans cette thèse, les termes compounding et connected seront utilisés de manière interchangeable.

fets économiques fondamentalement différents. A titre d'exemple, la sécheresse et les vagues de chaleur préconditionnent l'occurrence d'incendies de forêt (mentionné dans AghaKouchak et al., 2020). Un autre exemple est celui des tempêtes tropicales qui provoquent à la fois des dommages directs liés au vent et des dégâts supplémentaires par l'intermédiaire d'inondations et d'ondes de tempête, touchant particulièrement les communautés côtières.

Il existe notamment deux raisons pour lesquelles il est important de considérer les effets connectés dans les études économiques : premièrement, la réponse à de tels chocs peut être fondamentalement différente de celle provoquée par des dangers isolés. Par exemple, étudier l'effet des anomalies de température sur la production agricole, comme cela est couramment fait, peut être une bonne approche pour mesurer la fonction de réponse de la température sur la production des sols ou le rendement des cultures. Si, toutefois, l'anomalie de température se produit simultanément avec un événement de vague de chaleur[3], la perte de rendement observée ne peut être attribuée que partiellement à l'anomalie elle-même, mais également à la réduction de la productivité du travail de l'agriculteur chargé de la récolte. Un autre exemple serait l'estimation des périodes de reprise après un choc. Si l'on présume qu'un seul choc s'est produit, les implications politiques dérivées pour la reprise post-choc pourraient être erronées. Pour reprendre l'exemple précédent, on pourrait proposer un meilleur système d'irrigation en réponse à une anomalie de température, mais il faudrait également envisager des stratégies visant à réduire l'exposition à la chaleur des ouvriers agricoles.

Deuxièmement, si une étude économique est menée à l'échelle régionale, voire au niveau du district, la variation d'exposition à une diversité de dangers peut différer au sein d'une même zone géographique. Un village peut subir une inondation et une vague de chaleur, tandis qu'un autre village, situé à 30 km de distance mais dans le même district, pourrait connaître un glissement de terrain au lieu d'une inondation, en raison de fortes précipitations. Si ces autres événements sont exclus, cela peut conduire à l'omission de variables importantes dans l'estimation, causant un biais au niveau du résultat final. C'est particulièrement préoccupant pour les modèles en coupe transversale ou en panel incluant des variables climatiques (Hsiang, 2016). Il est donc important d'adopter une perspective multi-aléas, en particulier lorsqu'on mène des recherches ayant une pertinence politique.

[3] Les deux sont fortement liés, mais une vague de chaleur ajoute des facteurs environnementaux supplémentaires tels que l'humidité de l'air et la vitesse du vent.

Pour souligner davantage ce point, la Figure 2 illustre le nombre de types distincts de dangers (par exemple, inondation et vague de chaleur) enregistrés entre 2004 et 2015 à une échelle de 10x10 km.[4] Les données sont issues du Chapitre 2. La figure ne présente que onze types de dangers, notamment les sécheresses, les vagues de chaleur, les inondations, les séismes, les cyclones tropicaux, les vents extrêmes, les glissements de terrain, les tsunamis, les volcans et les incendies de forêt. La liste n'est pas exhaustive et pourrait être élargie à d'autres dimensions, mais elle suffit ici à des fins d'illustration.

Les deux graphiques montrent la part du PIB (à gauche) et de la population (à droite) exposée au nombre de classes de dangers[5]. À l'échelle mondiale, les graphiques révèlent que, pour les onze dangers représentés, moins de 7% du PIB et moins de 6% de la population n'ont été exposés à aucun événement extrême sur la période 2004–2015. Environ un quart de la population et du PIB a été exposé à un seul événement. Dans ces zones, il est possible de mener relativement facilement des études d'événement avec un seul type de traitement. Pour les 70% restants de l'échantillon, plus d'une classe de danger a été enregistrée, ce qui pourrait contaminer toute estimation économétrique si cela n'est pas pris en compte. Cela concerne notamment le sud-est des États-Unis, la zone sahélienne en Afrique, ainsi que certaines parties de l'Asie du Sud-Est, notamment le Vietnam, le Cambodge et le Laos. Ainsi, plus le nombre de classes de dangers de cette liste non exhaustive est élevé dans un même lieu, plus il est probable que les effets mesurés dans des études basées sur un seul événement soient en réalité influencés par des dangers omis ou des interactions entre plusieurs types de dangers.

Assez logiquement, cela interroge sur les raisons pour lesquelles la notion de chocs composés reste encore peu intégrée dans les travaux de recherche. Il y a à la fois des raisons pratiques et empiriques pour cela. Ce serait une erreur de supposer qu'aucune recherche n'est menée sur le sujet. Des travaux dans des domaines interdisciplinaires tels que ceux de Ridder et al. (2020) and Thalheimer, Choquette-Levy, and Garip (2022) ainsi que des travaux tenant compte des chocs cumulatifs (par exemple Di Falco et al., 2024) au fil du temps étudient la question des chocs connectés. Pourtant, en pratique, il y avait un manque de disponibilité de jeux de données homogénéisés qui permettent de démêler les différents chocs. Comme mentionné dans Botzen, Deschenes, and Sanders, 2019, de nombreuses études empiriques s'appuient sur la base de données

[4] La figure ne montre pas la fréquence à laquelle les dangers se produisent, mais uniquement le nombre de types différents observés entre 2004 et 2015.

[5] Les données sont basées sur les parts de PIB calculées à partir de Kummu, Kosonen, and Masoumzadeh Sayyar (2025) pour l'année 2015.

Fig. 2. Illustration de la présence des dangers et de leur importance relative

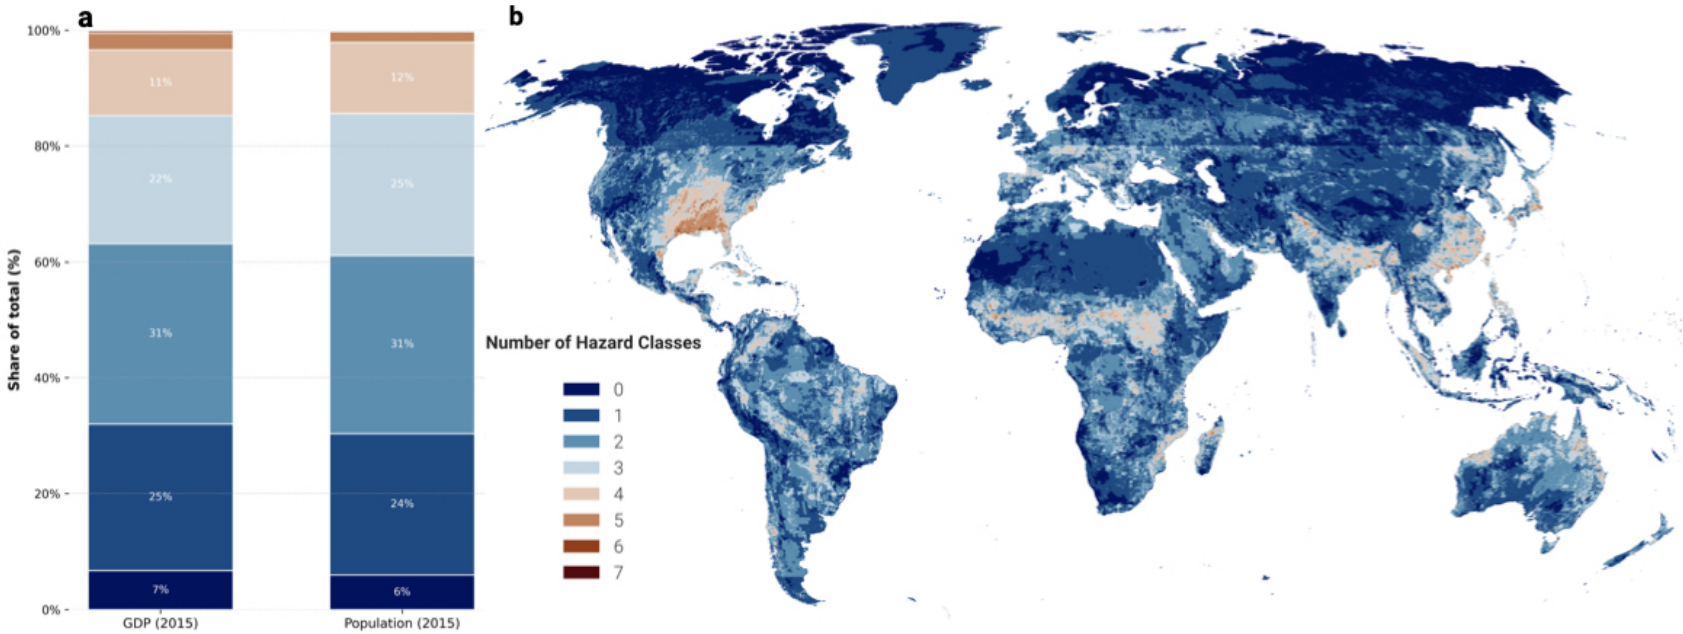

Note : Cette figure est basée sur l'ensemble de données compilé dans Reinhardt (2024a) portant sur un sous-ensemble de onze types de dangers courants (sécheresse, vague de chaleur, inondation, grêle, séismes, cyclone tropical, vents extrêmes, glissements de terrain, tsunami, volcan et incendies de forêt). La figure montre le nombre de classes de dangers distinctes enregistrées entre 2004 et 2015. Le panneau (a) calcule en outre la part du PIB mondial (à gauche) et de la population mondiale située dans des zones ayant le nombre respectif de classes de dangers distinctes. La population et le PIB sont tous deux calculés sur la base des valeurs de 2015. Les données sur le PIB proviennent de Kummu, Kosonen, and Masoumzadeh Sayyar (2025). Le panneau (b) divise l'estimation mondiale en sous-estimations régionales, indiquant la part du PIB et de la population selon le nombre de classes de dangers distinctes.

EM-DAT — décrite dans CRED and UC Louvain, 2024 — qui fournit des informations sur les catastrophes naturelles à l'échelle mondiale. Bien que largement utilisées, ces données ne précisent pas toujours la délimitation administrative exacte des événements, ce qui complique leur utilisation dans les études économiques à l'échelle régionale ou locale. En outre, des critiques de ces données incluent sa tendance à surestimer potentiellement les pays plus riches en raison de son accent sur les dommages économiques plutôt que sur l'intensité des aléas (Felbermayr and Gröschl, 2014). Bien que de nombreux jeux de données aient été proposés ces dernières années, ce n'est qu'avec la sortie récente de la base de données MYRIAD-HES (Claassen et al., 2023a) qu'un jeu de données complet pour les intensités locales des aléas est devenu accessible. Ces données sont compilées en définissant 11 types d'événements comme événements dangereux en fixant des seuils puis en compilant les données pour les années 2004 à 2017 au niveau journalier. Le chapitre 2 de cette thèse propose un nouveau jeu de données puisque, pendant son développement, il n'y avait pas de jeu de données standardisé qui incluait de manière fiable les intensités des aléas. La "Platform for Economic Analysis of Climate Hazards" (PEACH), un grand parti de cette these, fournit un jeu de données comparable avec des données d'entrée légèrement différentes à l'échelle mondiale mais avec une résolution mensuelle. Chaque fois que possible, elle s'abstient d'appliquer tout seuil, permettant aux chercheurs d'utiliser les intensités mensuelles pour calculer les dommages et les fonctions de réponse.

Au-delà des contraintes pratiques — telles que la disponibilité des données — plusieurs raisons empiriques expliquent pourquoi peu d'articles adoptent une perspective multi-aléas. Premièrement, les mêmes problèmes économétriques standards s'appliquent lors de l'utilisation d'aléas composés comme pour tout événement climatique. Cela inclut la non-linéarité des effets, les décalages temporels — comme les effets différés ou persistants —, les décalages spatiaux — lorsque les impacts s'étendent au-delà de la zone directement touchée —, ainsi que les défis liés à la modélisation correcte de la structure du terme d'erreur (Hsiang, 2016). Cependant, un autre défi est de savoir comment prendre correctement en compte l'interaction de différents types d'aléas, de facteurs environnementaux, ou de moteurs d'impact climatique. Un bon point de départ pourrait être de considérer les interactions entre les intensités de deux événements. Cependant, bien que cela soit relativement simple avec deux facteurs, cela devient de plus en plus difficile lorsque des types d'événements supplémentaires se cumulent. Avec deux intensités, par exemple, température et précipitations, et un résultat, on obtient deux réponses spécifiques aux aléas et un terme d'interaction de haute di-

mension. La Figure 3 permet de visualiser cela en prenant le cas du Nigéria et en combinant les effets de la température et des précipitations sur la santé des plantes au Nigeria. La santé des plantes est mesurée par le Indice de végétation par différence normalisée (IVDN). Le panneau central montre les effets individuels de la température et des précipitations sur la santé des plantes, tandis que le panneau de droite montre leurs interactions. La figure démontre que la prise en compte des interactions modifie les effets estimés des facteurs individuels. Si une troisième composante était ajoutée, comme la vitesse du vent, il deviendrait difficile pour les chercheurs de déterminer quelle combinaison entraîne les changements observés.

Une solution pourrait être l'utilisation d'une analyse de la sensibilité d'un système spécifique, comme un « test de résistance », où il est analysé à un moment spécifique et à un endroit spécifique, quels facteurs contribuent le plus aux changements dans les résultats souhaités (cf. Zscheischler, Martius, et al., 2020). Deuxièmement, différentes classes d'aléas pourraient être corrélées soit par une cause sous-jacente commune (dans le cas de téléconnexions telles que le El Niño – Oscillation australe ou le Oscillation de Madden-Julian), soit directement liées physiquement comme dans le cas de la pollution de l'air et des feux de forêt (voir par exemple Burke, Driscoll, et al., 2021). Si les classes d'aléas corrélées sont exclues des modèles économétriques, elles peuvent provoquer un biais de variable omise. De plus, comme mentionné par Hsiang (2016), ces variables non observées sont probablement autocorrélées dans le temps, ce qui biaise les erreurs standards de la régression. Même si elles sont incluses, elles peuvent limiter la possibilité d'inférences précises dans les modèles de régression paramétrique, potentiellement en raison de problèmes tels que la multicolinéarité imparfaite ou la modélisation inadéquate de leurs effets conjoints. Dans la littérature scientifique, dans de tels cas de dépendance entre variables, des méthodes statistiques, telles que les fonctions copules, ont été utilisées pour mieux prendre en compte l'interdépendance des variables (Zscheischler, Martius, et al., 2020). Enfin, se pose la question pratique du choix d'un modèle économétrique approprié. Les modèles paramétriques standards échouent souvent à prendre en compte les interactions complexes et non linéaires entre plusieurs aléas comme celles montrées dans la Figure 3. Les événements extrêmes, souvent traités comme des expériences naturelles singulières, nécessitent l'utilisation de modèles causals pour identifier les effets de manière robuste. Les modèles de discontinuité de régression (RD), en particulier dans l'espace, ne parviennent pas à tenir compte de la dépendance spatiale non observée au niveau local. Les régions voisines diffèrent souvent dans leur exposition au risque

climatique, leurs trajectoires de croissance économique et leurs stratégies d'adaptation climatique, ce qui complique l'utilisation des modèles RD spatiaux. Les essais contrôlés randomisés ne sont pas applicables étant donné que le traitement des événements extrêmes peut être aléatoire dans le temps mais pas dans l'espace. Les méthodes de contrôle synthétique rencontrent quant a elles des défis dans le contexte des événements extrêmes en raison des effets potentiels d'équilibre général et des retombées qui peuvent biaiser les estimations même avec des groupes de contrôle soigneusement construits. Enfin, et cette méthode a été utilisée largement tout au long de la thèse, il existe les modèles de différence-en-différence. Ces modèles permettent de capturer l'effet en comparant les zones exposées avec des zones jamais traitées ou pas encore traitées. Un défi clé est d'assurer la validité de l'hypothèse de tendances parallèles, en particulier lorsque plusieurs aléas sont présents ou lorsque des événements extrêmes réapparaissent durant la période d'étude. Pour tenir compte de ces complexités, tous les chapitres de la thèse s'appuient sous une forme ou une autre sur des estimateurs récents de la littérature sur la différence-en-différence. Bien que non parfaits, certains de ces nouveaux estimateurs permettent la réversibilité du traitement, ce qui est essentiel dans les études sur les catastrophes.

Dans le paragraphe suivant, je décris comment cette thèse définit les événements connectés et comment elle aborde les préoccupations méthodologiques mentionnées ci-dessus. Les chapitres 2 à 4 incluent la notion d'événements composés ou connectés (appelés événements « complexes » dans le Chapitre 3). J'adopte une approche simplifiée en définissant un événement composé comme un événement qui se produit dans la même zone (« événement composé spatial ») ou dans une certaine période de temps (« événement cumulatif temporel »). Dans le Chapitre 4, j'explore également différentes combinaisons de séquences d'événements.

Pour tenir compte des interdépendances entre les types d'aléas, les Chapitres 3 et 4 s'appuient sur des termes d'interaction. Bien que des fonctions de réponse non linéaires, comme montré et décrit ci-dessus, seraient préférables, j'opte pour une approche alternative en raison d'une hétérogénéité substantielle dans les formats des données utilisées. Au lieu de s'appuyer sur des mesures continues d'intensité, les événements individuels sont transformés en variables muettes de traitement binaires basées sur des seuils statistiques (« upper tail events ») ou issus de la littérature. Pour certains aléas, comme les anomalies de température, les "upper tail events" sont définis par rapport aux tendances historiques spécifiques à chaque lieu. Pour d'autres aléas, tels que les cyclones tropicaux, un seuil global est utilisé. Bien que les mesures con-

Fig. 3. Exemple de non-linéarités dans la dimension des impacts

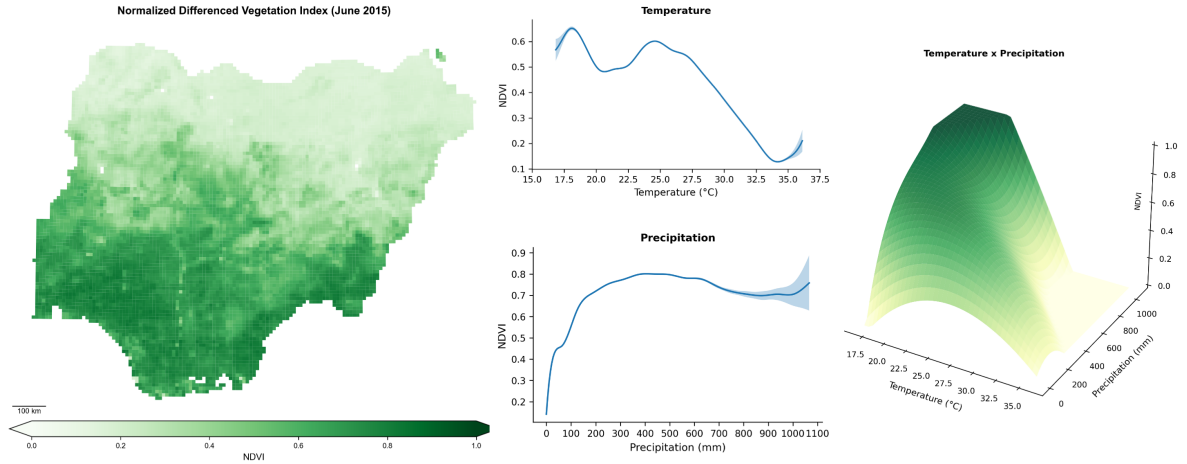

Note : Cette figure fournit un exemple des effets de la température et des précipitations sur la santé des plantes, mesurée par l'indice de végétation par différence normalisée (NDVI), au Nigeria. La figure de gauche illustre, à titre d'exemple, la distribution du NDVI au Nigeria pour un mois aléatoire, juin 2015. Les données NDVI sont issues de Li, Cao, Zhu, et al. (2023). Les deux figures centrales tracent une fonction de réponse de la température moyenne mensuelle et des précipitations moyennes en mm sur le NDVI, en utilisant un modèle additif généralisé (GAM) avec quatre polynômes. Les données proviennent d'ERA-5 (Muñoz Sabater, 2019) et de CHIRPS (Funk, Peterson, Landsfeld, Pedreros, Verdin, Shukla, et al., 2015). Le graphique de droite montre l'interaction entre les deux coefficients sur le NDVI. Les données dérivées sont basées sur le Chapitre 2 de cette thèse.

tinues des aléas permettraient l'estimation de fonctions dose-effet, un tel cadre n'est pas possible dans cette thèse en raison de l'hétérogénéité significative dans les formats de données inclus. Certains aléas sont rapportés de manière catégorielle, d'autres à travers des scores normalisés, et certains sous forme binaire. Cette incohérence entre les formats de données complique fortement la construction de paires d'interaction et limite la pertinence des modèles non linéaires. En contraste, l'exemple de la Figure 3 utilise la température et les précipitations, deux variables bien étudiées dans la littérature, qui ne sont pas représentatives de tous les aléas dans cette thèse. Les aléas, tels que les feux de forêt, par exemple, sont plus difficiles à intégrer dans un cadre non linéaire. Par conséquent, cette dissertation suit de près la littérature sur les études d'événements, qui repose généralement davantage sur des variables muettes de traitement binaires. Pour permettre une certaine forme de non-linéarité, les Chapitres 3 et 4 incluent et discutent des termes quadratiques dans les modèles paramétriques de base. Je reconnais que cela se fait au prix de l'omission d'une partie de l'hétérogénéité sous-jacente dans l'impact des événements, ce qui est discuté plus en détail dans les

chapitres correspondants. En conséquence, dans les Chapitres 3 et 4, je souligne également d'autres facteurs que les chercheurs pourraient vouloir prendre en compte pour évaluer si leurs résultats sont véritablement causaux. Une autre contribution de cette thèse est l'introduction d'un jeu de données inédit, PEACH, au Chapitre 2, qui répond au problème décrit de l'absence d'une base de données unifiée pour l'analyse des catastrophes. Les préoccupations liées à la corrélation temporelle et spatiale, ainsi qu'à la structure du terme d'erreur, sont abordées en appliquant différents niveaux d'agrégation temporelle et spatiale. Tous les modèles des Chapitres 2 à 4 incorporent la notion de "CID", mais la nomment dans les Chapitres 3 et 4 « événements extrêmes » ou « chocs climatiques » pour des raisons de compréhension. Le Chapitre 1 propose une perspective alternative sur la littérature. En étudiant les flux d'investissement, l'objectif de ce chapitre est de soutenir que les effets des événements extrêmes destructeurs vont au-delà de la seule destruction du capital physique. Les régions spatialement ou économiquement connectées à une zone touchée peuvent être affectées indirectement par le choc, en particulier lorsque les chocs surviennent en amont, c'est-à-dire sur le versant de l'offre de l'entreprise. Bien que le risque multi-aléas ne soit pas l'objectif principal de l'article, la base du modèle prend néanmoins en compte la présence d'autres événements extrêmes.

La contribution de cette thèse est donc triple : le Chapitre 1 met en évidence, à travers l'exemple de l'investissement étranger, que l'effet des événements extrêmes va au-delà de la seule destruction des actifs et peut se propager le long des chaînes d' approvisionnement jusqu'à des entreprises éloignées. Le Chapitre 2 présente un nouveau jeu de données en libre accès afin de mieux estimer les interactions des événements extrêmes à l'échelle locale. Enfin, les Chapitres 3 et 4 revisitent des sujets largement étudiés en économie du développement, en se concentrant sur les interactions entre différents types d'aléas et leurs implications pour les résultats. La section suivante propose un aperçu plus détaillé de l'ensemble des chapitres de la thèse.

Contribution détaillée de chaque chapitre de la thèse

Les paragraphes suivants présentent chacun des chapitres et fournissent une vue d'ensemble. Ils soulignent leur contribution à la littérature sur les événements extrêmes et connexes dans les pays du Sud.

Chapitre 1

Le premier chapitre est le seul chapitre qui ne se concentre pas particulièrement sur les chocs connectés ou les chocs cumulatifs. Au lieu de cela, l'article examine l'impact des événements extrêmes destructeurs, en particulier les tremblements de terre, sur les flux d'investissements étrangers à l'intérieur d'un pays. Comme pays d'étude, cet article se concentre sur l'Indonésie en raison de son exposition fréquente à une large gamme d'événements extrêmes. La position de l'Indonésie à proximité de la Ceinture de Feu du Pacifique rend les tremblements de terre puissants et largement répartis à travers le pays. La question de savoir comment les événements extrêmes affectent l'investissement a été largement étudiée dans la littérature, mais la plupart de ces articles se concentrent sur des dynamiques au niveau national. À ma connaissance, il n'existe que très peu d'études qui incluent des effets des chocs à l'intérieur d'un pays. Un exemple notable est l'exemple de Friedt and Toner-Rodgers (2022) pour les inondations en Inde. Leur article est le plus proche de mon article. Au-delà de cela, plusieurs études se penchent sur la dynamique des entreprises dans les périodes qui suivent ces événements. Un article étroitement lié ici est celui de Carvalho et al. (2021) qui étudie les effets de la chaîne d'approvisionnement des entreprises au Japon dans les suites d'un tremblement de terre particulièrement fort. L'article contribue à la littérature en se concentrant, de manière similaire à Carvalho et al. (2021), sur les effets de la chaîne d'approvisionnement des tremblements de terre, mais sans utiliser de données au niveau des entreprises, comme dans Friedt and Toner-Rodgers (2022). De plus, j'étudie les retombées spatiales et, dans l'annexe de l'article, les effets en tenant compte d'autres aléas.

Pour cela, l'article combine des données sur les investissements directs étrangers entrants provenant de 413 districts indonésiens entre 2003 et 2019. Ces données sont maintenues par le Conseil de Coordination de l'Investissement Indonésien et sont fournies au niveau sectoriel et mensuel. Pour obtenir une approximation de l'intensité des tremblements de terre, les données sur les investissements étrangers sont combinées avec 896 cartes géophysiques détaillées qui illustrent l'empreinte des événements sismiques. Les données sont prises de l'Institut d'Études Géologiques des États-Unis et sont appelées « ShakeMaps ». L'intensité est basée sur des modèles géophysiques qui estiment l'intensité des secousses à partir de données d'observation in situ provenant de stations d'enregistrement. Contrairement aux études précédentes, l'article applique deux conditions qui doivent être remplies pour qu'un tremblement de terre soit considéré comme significatif. Premièrement, le marché du travail local

doit être perturbé, ce qui signifie que des personnes dans le district ont été enregistrées comme affectées par un tremblement de terre. Deuxièmement, le tremblement de terre doit avoir causé au moins des dommages mineurs. Pour cela, j'applique un seuil établi précédemment dans la littérature, qui est en outre validé dans la section de robustesse. Bien que l'information sur l'emplacement des tremblements de terre et des investissements étrangers soit incluse, il manque l'information sur la façon dont l'activité économique est reliée entre les provinces. À cette fin, nous renforçons le modèle avec une table entrée-sortie multirégionale de 2016 pour relier l'activité économique sectorielle à travers l'espace à l'intérieur du pays. Ainsi, je peux tester si les chocs sismiques se propagent dans l'économie ou restent confinés aux zones touchées.

Méthodologiquement, j'applique un estimateur basé sur une étude d'événements utilisant un modèle à effets fixes bidirectionnels. Cependant, suivant la discussion de Goodman-Bacon (2021), les estimations dans un tel modèle peuvent être biaisées si les tremblements de terre ne se produisent pas au même moment, ce qui est le cas dans cet article. Ainsi, dans la section de robustesse, j'applique une correction pour cela et je compare le résultat de base avec des estimateurs alternatifs issus de la littérature récente sur la différence de différences. J'observe que les flux d'investissement sont temporairement réduits de presque 90% l'année suivant un tremblement de terre. Toutefois, toutes les zones ne sont pas affecté de la même manière. Les effets sont principalement concentrées dans les zones avec une exposition historique au risque plus faible et un stock d'investissement étranger plus élevé. Cela suggère que l'anticipation du risque, ainsi que le nombre total d'entreprises dans lesquelles la firme étrangère a investi, peuvent jouer un rôle dans la réponse à un choc. Ensuite, je passe aux effets indirects du tremblement de terre. D'abord, en considérant les retombées spatiales — bien que ce ne soit pas le cœur de cet article — j'observe que les retombées en provenance des districts voisins ont peu d'effet sur les résultats principaux. Cela s'explique probablement par la forte concentration de l'investissement étranger autour des centres économiques et des zones économiques spéciales. En revanche, la réduction de l'investissement étranger peut être partiellement attribuée à des chocs induits par l'offre. Si les fournisseurs d'une entreprise étrangère sont impactés par un tremblement de terre, surtout dans la fabrication, cela tend à réduire significativement l'entrée étrangère. Du côté de la demande, le secteur de la construction semble particulièrement bénéficier lorsque les clients sont affectés par des tremblements de terre. Dans l'ensemble, l'étude met en évidence le cas d'un seul type d'aléa et d'un seul pays,

montrant que, surtout pour l'analyse des impacts d'aléas à l'échelle locale, il est également important de considérer les effets indirects comme les retombées spatiales ou de réseau.

Chapitre 2

Le deuxième chapitre tente de combler le fossé mentionné ci-dessus concernant l'absence d'un cadre de données unifié pour étudier la co-occurrence des événements extrêmes. La plupart des articles dans la littérature s'appuient soit sur des aléas individuels, soit sur des bases de données d'aléas compilées telles que EM-DAT (CRED and UC Louvain, 2024). Ce n'est que récemment qu'il y a eu des efforts pour créer une base de données mondiale à haute résolution des événements extrêmes. En ce qui me concerne, je souhaitais disposer d'un estimateur davantage conçu pour des séries temporelles par localisation, tout en tenant compte des conditions socio-démographiques et économiques propres à chaque zone. Ainsi, j'ai développé la Plateforme d'Analyse Économique des Risques Climatiques (PEACH) qui fournit un ensemble de données global, quadrillé, compilé et homogénéisé pour 14 facteurs d'impact climatique et 3 aléas non climatiques. Ceux-ci vont des aléas standards tels que les anomalies de température ou de précipitations, aux grêles, feux de forêt ou cyclones tropicaux. Les aléas non influencés par le climat incluent notamment les tsunamis, les tremblements de terre et la pollution de l'air. L'ensemble de données combine ce grand ensemble de variables pour chaque mois de 2000 à 2019, mais n'atteint une couverture complète pour les 17 types d'aléas qu'entre les années 2004 et 2015, en raison d'une disponibilité limitée des données sources.

La résolution spatiale est d'environ 10x10 km à l'échelle mondiale avec deux limitations mineures : pour les précipitations, je n'ai qu'une couverture quasi-globale, et une partie de l'hémisphère nord est manquante. Deuxièmement, étant donné que j'utilise une projection de Mercator pour les données incluses afin de suivre les standards de cartographie, cela peut entraîner des distorsions dans les fonctions non linéaires en fonction de la distance à l'équateur, ce que les lecteurs sont invités à prendre en compte. Pour plusieurs des 17 aléas mentionnés, des valeurs d'intensité sont disponibles, cependant, pour certains aléas, seules des variables binaires ou des variables de comptage sont incluses. Contrairement aux études précédentes, pour les dimensions les plus importantes, telles que la sécheresse, les vagues de chaleur, la température, les précipitations et la pollution de l'air, je m'abstiens d'imposer un seuil, lorsque cela est

possible, afin de permettre aux chercheurs de choisir leur propre intensité respective. Cela doit permettre une meilleure calibration des fonctions de dommages des résultats économiques en fonction de l'intensité de l'aléa. L'objectif principal de cet ensemble de données est de faciliter l'accès à des données qui exigent généralement de nombreuses compétences en programmation pour être converties en un format de recherche utilisable en sciences sociales. Il fournit de meilleures informations sur la distribution mondiale des événements extrêmes sur deux décennies et doit aider les chercheurs à étudier l'impact des événements connectés ou cumulatifs sur les résultats économiques.

Comme les chercheurs en sciences sociales peuvent s'intéresser aux résultats économiques, j'ajoute également un ensemble d'indicateurs socio-économiques précompilés afin de faciliter l'utilisation des données. Cela inclut les lumières nocturnes, la population, des informations sur la part des infrastructures critiques et les données d'utilisation des sols. Ce chapitre contribue donc à la littérature en fournissant aux chercheurs un ensemble de données supplémentaire, facile à utiliser pour l'étude, venant compléter les ensembles de données préexistants. Les données sont validées par rapport aux données des stations météorologiques et par rapport à la base de données EM-DAT et fournissent une qualité de correspondance suffisante. Le plan futur inclut l'offre du code complet de réplication exécutable sous la forme d'un package Python afin de permettre aux chercheurs d'incorporer des catégories d'aléas supplémentaires encore non couvertes.

Chapitre 3

Le troisième chapitre intègre le nouveau jeu de données du chapitre 2 et revisite la littérature bien établie sur les changements de population et la migration en Afrique. La plupart de ces articles se concentrent soit au niveau des ménages, soit sur les effets de la migration internationale en utilisant des estimations de migration au niveau national (Cattaneo, Beine, et al., 2019). Cet article recentre l'analyse sur une perspective macroéconomique et évite ainsi d'attribuer l'effet observé uniquement à la migration. Contrairement à la littérature à laquelle cet article est opposé, cet article ne se concentre pas sur un ou deux événements extrêmes, mais prend en compte, en accord avec l'objectif de cette thèse, neuf aléas communs à travers l'Afrique subsaharienne. L'article répond à la question de savoir comment ces événements extrêmes influencent le changement net de population dans le contexte du continent. La variable de résultat

d'intérêt est le changement net de population entre 2000 et 2019 provenant de Niva et al. (2023) et mesure les changements nets de population interpolés sur une base annuelle et à travers 3863 districts en Afrique. Je m'abstiens d'appeler les changements de population "migration" même si une grande partie du changement net de population est probablement entraînée par des tendances migratoires. Les données sur les événements extrêmes sont basées sur le deuxième chapitre de cette thèse. J'inclus des informations sur les événements extrêmes concernant les anomalies de température, les chocs de précipitations, les vagues de chaleur, les sécheresses, les tempêtes, les inondations, la grêle, les incendies de forêt et les invasions de criquets. Tous ont été sélectionnés car ils affectent soit l'agriculture soit pourraient avoir d'autres impacts directs sur les moyens de subsistance des populations. Un problème clé est que les données dérivées du chapitre 2 utilisent principalement des intensités, mais certains aléas ne disposent que de variables indicatrices binaires ou de variables catégorielles. De plus, elles sont disponibles à une résolution de 10 km pour chaque mois. Le résultat, en revanche, a une résolution annuelle au niveau des districts. Ainsi, comme mentionné précédemment, je convertis les intensités en variables binaires en appliquant des seuils spécifiques à chaque lieu qui tiennent compte des tendances historiques, chaque fois que cela est possible. Pour certains aléas, nous nous basons sur des seuils issus de la littérature. En complément, nous appliquons un seuil de surface pour tenir compte de l'exposition intra-district. Autrement dit, un district n'est considéré comme traité que si une proportion minimale de sa superficie a été affectée par un choc climatique. Ensuite, ces indicateurs binaires mensuels au niveau des districts sont agrégés et additionnés à un niveau annuel.

L'article étudie les effets de ces événements extrêmes sur le changement net de population en Afrique. L'analyse est ensuite décomposés en deux parties. Premièrement, je suis une approche plus standard de la question en examinant quels sont les effets des événements extrêmes individuels sur le changement net de population. J'utilise ici de manière simplifiée une variable muette binaire de traitement qui indique si un aléa spécifique a été observé dans le district ou non. En revanche, contrairement à la littérature précédente, je prends en compte les aléas simultanés, c'est-à-dire la présence d'autres événements extrêmes se produisant au même moment, afin d'éviter que l'effet observé ne soit attribuable à des facteurs non observés. Ensuite, dans la deuxième partie, je divise la variable de traitement des aléas en trois régimes de traitement différents. Un événement enregistré au cours de l'année peut appartenir à l'un des trois régimes suivants : un événement simple est défini lorsque seul un type d'aléa est enregistré au

cours d'une année donnée. Un exemple serait une inondation unique qui affecte un district en 2005. Le deuxième type concerne les soi-disant événements cumulatifs. Ces événements font référence à une situation où le type d'aléa se produit plus d'une fois au cours d'une année. Un exemple serait des tempêtes tropicales qui surviennent deux fois au cours d'une année, mais à des mois différents, affectant le même district en raison de déplacements de la zone de convergence intertropicale. Le troisième type est un type nouveau basé sur la littérature en sciences climatiques et appelé événements composés. Un événement composé est défini dans cet article comme un chevauchement temporel de deux types d'événements dans les mêmes mois au sein du même district. Un exemple serait une vague de chaleur documentée dans un district qui entraîne une évaporation d'air chaud, laquelle provoque à son tour de la convection et des précipitations excessives causant une inondation dans ce district. Je discute en détail dans l'annexe de l'article la manière dont cet événement est construit et quelles en sont les limites. J'inclus ainsi dans la version finale les interactions temporellement superposées des 38 combinaisons différentes. Comme cela pourrait entraîner des problèmes de multicolinéarité imparfaite, je mène plusieurs tests de robustesse, y compris un modèle PDS-Lasso, afin de valider que mes résultats de référence ne sont pas biaisés. Les deux parties, la première et la seconde, reposent largement sur un modèle de panel avec effets fixes pour le district, l'année et la région continentale par année. Les tests de robustesse sont principalement effectués sur la première partie de l'article et je ne décris que les tendances des effets les plus significatifs dans la deuxième partie, compte tenu du nombre de coefficients. Dans le paragraphe suivant, je vais décrire ce que nous trouvons dans ce chapitre.

Initialement, dans le scénario de référence, en se concentrant sur les aléas uniques avec une variable muette annuelle par aléa, je constate que les inondations, en particulier, réduisent le changement net de population, entraînant notamment une émigration. La grêle, en revanche, tend à augmenter l'immigration. Les résultats sont relativement robustes selon la spécification, bien que certains effets soient masqués en raison de l'inclusion des effets fixes région continentale-année. Ensuite, en passant à la deuxième partie de l'article, le premier résultat principal est présenté. Alors que de nombreux coefficients d'aléas individuels sont non significatifs, la prise en compte des paires d'interaction d'événements extrêmes permet d'obtenir une image fondamentalement plus riche des effets des événements extrêmes. De nombreux aléas auparavant non significatifs deviennent des facteurs significatifs de changement de population, notamment en combinaison avec d'autres événements. Pour certains, notamment

les inondations et les sécheresses, l'accumulation du même événement au cours de l'année conduit également à des réponses plus fortes. Ainsi, les modèles simples ne tenant compte que des traitements binaires risquent de ne pas capturer adéquatement la complexité des impacts des événements extrêmes. Les chercheurs devraient plutôt considérer les paires d'aléas individuelles comme des « nouveaux » événements.

Deuxièmement, de manière similaire aux études de cette littérature, l'article teste une hétérogénéité potentielle ainsi que les mécanismes qui pourraient expliquer les effets documentés. La plupart des effets précédemment documentés sont observés dans des districts plus ruraux, plus pauvres et plus vulnérables au climat. Les personnes dans ces districts pourraient ne pas avoir d'alternative autre que de migrer en raison du manque d'options, ou bien de s'abstenir de migrer si les ressources financières suffisantes ne sont pas présentes au sein du ménage. En outre, je documente que certaines combinaisons d'aléas ont des effets marqués sur les indices de sécurité alimentaire ainsi que sur la santé et la production des cultures. Toutefois, il existe une grande hétérogénéité. Par exemple, je constate que la réponse des aléas sur le maïs est généralement plus faible que sur le manioc, ce qui révèle que, au-delà de la simple présence de l'aléa, la composition des cultures d'une zone pourrait également jouer un rôle dans la réponse à un choc. Je constate que les sécheresses et les événements liés à la chaleur semblent particulièrement être des facteurs majeurs de pertes en production agricole.

Troisièmement, les tests de robustesse de cet article révèlent que la méthodologie proposée souffre en partie de biais liés à des non-linéarités cachées, à des effets de traitement dynamiques ainsi qu'à des effets de débordement spatial. Cela révèle ainsi une complexité structurelle plus profonde de la manière dont les événements extrêmes affectent les résultats socioéconomiques. De nombreux articles dans la littérature se concentrent soit sur les effets de traitement dynamiques de chocs singuliers, soit sur des modèles permettant des effets non linéaires des aléas. Cependant, ne pas tenir compte des trois types de biais potentiels risque de produire des résultats erronés. En outre, comme mentionné précédemment, il peut être pertinent de tenir compte de la présence d'autres aléas et d'être prudent avant d'attribuer directement une causalité aux effets observés. Par conséquent, ce chapitre de la thèse contribue à la littérature en tenant compte activement des types d'interactions d'aléas au cours de l'année, en revisitant les mécanismes potentiels qui peuvent expliquer pourquoi les événements extrêmes influencent les changements de population, et en soulignant les risques que les chercheurs doivent considérer lorsqu'ils s'appuient sur des variables muettes d'aléas.

Chapitre 4

Le dernier chapitre ne s'intéresse plus à la population rurale mais examine les effets des aléas sur les destinations potentielles des populations déplacées, en particulier les zones urbaines. L'article s'interroge ainsi sur les effets à court terme de quatre des phénomènes météorologiques les plus courants. Nous passons ici d'une qualification d'événements extrêmes à celle de chocs météorologiques spécifiques, afin de souligner la dimension temporelle plutôt brève de leur impact et l'intensité de celui-ci, influencée par les conditions météorologiques locales. Les chocs météorologiques considérés sont les inondations, les vagues de chaleur, les tempêtes et les sécheresses, en raison de leur forte probabilité d'être affectés par le changement climatique. LLa question de recherche qui en découle porte donc sur les effets à court terme que ces quatre aléas exercent sur l'expansion urbaine horizontale physique[6] dans un échantillon de 5 721 villes d'Afrique subsaharienne. L'expansion urbaine horizontale physique se concentre notamment sur la croissance spatiale des villes, plutôt que sur leur croissance en hauteur ou en densité. Les études antérieures se sont généralement penchées sur la dimension d'impact de certains aléas, le plus souvent sur la manière dont les sécheresses ou les anomalies de précipitations influencent les tendances à l'urbanisation. Cependant, en accord avec cette thèse, j'ajoute les chocs simultanément afin de capter les effets de contagion par d'autres événements extrêmes. De plus, au lieu de considérer les aléas de manière individuelle ou binaire, ceux-ci sont ici ajoutés comme somme totale des chocs. Les chocs sont définis, comme dans le deuxième article, comme des séquences consécutives de mois durant lesquelles des seuils prédéfinis sont dépassés, ces seuils étant établis sur des bases historiques. Le modèle final comporte ainsi quatre variables cumulées, représentant le nombre de chocs de chaque type d'aléa. Ensuite, j'applique un modèle de panel sur 20 ans, en contrôlant les effets fixes de la ville, de l'année, et de la région continentale par année. Les effets fixes région-années visent à prendre en compte les conditions atmosphériques à grande échelle pouvant influencer la probabilité de survenue des aléas. Un exemple notable serait l'Oscillation de Madden-Julian, qui pourrait influencer les régimes de précipitations régionales. Le résultat d'intérêt est une mesure appelée « surface imperméable ». En bref, elle mesure la surface d'une ville recouverte de béton. Cela a deux implications majeures pour cette étude : cela entrave souvent l'écoulement de l'eau, augmentant ainsi le risque

[6] L'expansion urbaine horizontale physique désigne l'extension spatiale du bâti et du infrastructure en périphérie des villes, en opposition à la densification verticale ou à la réhabilitation du tissu urbain existant.

d'inondation, et cela réfléchit la chaleur, contribuant ainsi aux vagues de chaleur en milieu urbain. Ainsi, au-delà de son intérêt pratique, il s'agit d'un indicateur important de croissance urbaine durable. Ces données sont disponibles à une résolution annuelle de 30 mètres et sont converties en superficie totale avec surfaces imperméables dans les limites urbaines fixées de 2020, c'est-à-dire à la fin de la période étudiée. Pour valider ces données, elles sont croisées avec d'autres mesures courantes de l'expansion urbaine.

En combinant les points précédents — à savoir la somme des quatre aléas, le modèle de régression en panel avec effets fixes régionaux, et la mesure d'expansion urbaine — que révèle l'article ? Je trouve des effets forts et persistants des inondations, qui réduisent l'expansion urbaine entre 3 et 9%. Les autres aléas montrent à peine des effets significatifs. Les sécheresses semblent légèrement accélérer l'expansion urbaine, mais les effets ne sont que faiblement significatifs. Une analyse plus approfondie révèle que les effets des inondations concernent notamment les villes plus pauvres et plus rurales d'Afrique de l'Ouest. Contrairement aux études précédentes, je ne constate pas que seules les villes primaires des pays croissent ; une grande partie des effets de croissance se concentre en réalité dans des villes de taille moyenne. Étant donné l'ampleur des effets significatifs des inondations, une attention particulière est accordée aux types d'inondations susceptibles d'en être la cause. En m'appuyant sur des données d'anomalies des débits fluviaux et des précipitations, je conclus que les inondations pluviales — donc provoquées par la pluie — sont probablement le principal facteur conduisant à une réduction de la croissance urbaine.

Ensuite, je propose trois extensions au modèle de base : premièrement — et il s'agit de l'extension principale — j'introduis des interactions entre les quatre aléas, conformément à l'idée de la combinaison des risques (« hazard compounding »). Au lieu de me limiter aux interactions ou chevauchements de deux aléas survenant le même mois, j'examine également si l'ordre des événements joue un rôle. Je constate, entre autres, que les vagues de chaleur précédant des inondations semblent avoir un effet significatif sur l'expansion urbaine. Les vagues de chaleur, prises isolément, ont peu d'effet. En les intéragissant avec les innondations, on observe, cependant, davantage d'effets. Les tempêtes, individuellement, montrent des effets faibles, probablement parce qu'elles n'affectent qu'une petite fraction des villes africaines situées sur la côte est. Bien que des effets positifs significatifs sur la croissance soient observés pour les villes plus riches, la faible taille de l'échantillon de villes concernées exige une certaine prudence quant à la généralisation des résultats. Des effets faibles sont également iden-

tifiés pour les interactions tempête-inondation, une combinaison fréquente à la suite de tempêtes tropicales.

Ensuite, j'examine si les effets varient selon que l'analyse soit menée à l'échelle de la ville (en intensité) ou à celle du bassin versant urbain (en extension). J'explore ici si les effets changent en fonction de la superficie environnante affectée. Cela semble être le cas pour les sécheresses. Les sécheresses au niveau de la ville ont peu d'importance. Toutefois, lorsqu'on considère la zone de captation urbaine, les sécheresses entraînent une accélération significative de la croissance urbaine d'environ 3%. Ce résultat est cohérent avec la littérature existante.

Troisièmement et pour finir, la plupart des articles sur l'expansion urbaine ne prennent pas en compte l'historique des aléas. Pour contribuer à cette littérature, je m'appuie sur un estimateur d'événement récent proposé par Imai, Kim, and Wang (2023) et je l'applique à l'analyse. En substance, cet estimateur suit la logique des études d'événement et construit des groupes de contrôle à partir d'un ensemble de variables de contrôle et d'historiques de traitement similaires. En appliquant cet estimateur, je découvre que si les sécheresses semblent initialement accélérer l'expansion urbaine, la prise en compte de l'historique des sécheresses inverse cet effet : une sécheresse supplémentaire conduit alors à une réduction de la croissance urbaine.

Face à tous ces effets sur la croissance urbaine, l'article discute également de la question suivante : une réduction de l'expansion urbaine doit-elle être considérée comme un résultat positif ? En soi, la croissance urbaine est neutre : elle peut permettre aux villes de mieux s'adapter aux événements. Mais si elle conduit à des installations informelles, elle peut aussi exposer davantage de personnes aux risques climatiques. Ainsi, les effets négatifs des inondations sur la croissance s'accompagnent d'une intensification de la population urbaine, d'une augmentation des quartiers informels, et d'une plus grande part de la population locale exposée au risque d'inondation futur. Je documente également une probabilité accrue de troubles violents à la suite de tels événements. La réduction de l'expansion urbaine est donc liée à des conséquences indésirables.

En conclusion, ce chapitre complète la littérature existante en montrant, comme le chapitre 3, que certaines combinaisons d'événements peuvent expliquer une partie des effets observés initialement. Les effets des chocs météorologiques sur l'expansion urbaine dépendent du fait que les chercheurs tiennent compte ou non de l'historique des traitements, de l'échelle spatiale considérée et des interactions entre les différents chocs météorologiques. D'un point de vue politique, ce chapitre met en évidence que

les inondations semblent exercer une pression accrue sur les villes les plus pauvres, qui connaissent une croissance plus importante des quartiers informels et une possible détérioration des conditions de vie. De manière générale, l'article s'inscrit pleinement dans l'objectif de la thèse, qui vise à mieux comprendre les effets des événements composés dans les pays en développement ainsi que les défis empiriques que posent ce type d'études.

Conclusion

Ensemble, ces quatre chapitres analysent la complexité que les événements extrêmes imposent aux résultats économiques. Ils mettent en évidence deux défis majeurs : la difficulté de modéliser de tels événements et la complexité des mécanismes sous-jacents—en particulier dans les pays à revenu faible et intermédiaire. Cette thèse apporte trois contributions principales à la littérature en économie du développement. Premièrement, elle introduit un nouveau jeu de données accessible, permettant aux chercheurs—en particulier ceux moins familiers avec la géoscience ou la télédétection—d'étudier les événements extrêmes, en mettant l'accent sur les environnements multirisques. Deuxièmement, elle offre de nouvelles perspectives sur l'impact des événements composés (définis comme des événements de types différents survenant au cours d'un même mois) et cumulatifs sur l'évolution de la population et l'urbanisation. Ces résultats suggèrent que les effets couramment observés peuvent être largement influencés par les interactions composées des événements extrêmes. Troisièmement, la thèse cherche à combler un fossé interdisciplinaire entre les sciences climatiques et l'économie en appliquant des modèles économiques empiriques à des données issues de la météorologie, des sciences du climat et de la géographie.

Les résultats empiriques confirment ces observations. Le chapitre 1 montre que les aléas ne se limitent pas à détruire les infrastructures physiques ou à mettre en danger les populations affectées, mais qu'ils exercent également des effets systémiques plus larges—perturbant les chaînes d'approvisionnement et se répercutant sur les régions voisines. Cela est illustré à travers le cas des flux d'investissements étrangers en Indonésie. Les chapitres 3 et 4 déplacent l'analyse géographique vers l'Afrique, révélant que les régions rurales, vulnérables et pauvres subissent des impacts démographiques significatifs à la suite d'événements composés. Les zones urbaines, quant à elles, présentent des trajectoires de croissance plus lentes, en particulier après des interactions entre inondations et vagues de chaleur. Toutefois, les deux chapitres soulignent également une hétérogénéité importante des effets observés, influencée par des facteurs tels que le type d'aléa, la richesse régionale, l'échelle spatiale et la dimension

spécifique de l'impact étudiée.

Bien que le jeu de données et les stratégies empiriques développés dans cette thèse visent à améliorer notre compréhension des conséquences économiques des événements extrêmes, ils présentent plusieurs limites. Premièrement, comme discuté à divers endroits de la thèse, tous les chapitres s'efforcent de contrôler autant que possible les menaces pesant sur l'identification. Toutefois, les modèles empiriques présentés ne reflètent pas une relation parfaitement causale. Identifier des effets causaux dans un contexte de chocs qui se chevauchent, d'interdépendance spatiale et de réponses dynamiques différées constitue un défi complexe. Par exemple, estimer l'effet causal d'événements composés sur la migration requiert une compréhension des dimensions par lesquelles les deux événements extrêmes affectent la migration, ainsi que de leur relation temporelle, spatiale et en termes d'intensité-impact. Cela exclut déjà la question de l'adaptation, où différents individus ou communautés peuvent réagir de manière différenciée en fonction de leur histoire d'exposition ou de l'apprentissage par les réseaux sociaux. Tous ces facteurs rendent difficile une identification propre et strictement causale.

Deuxièmement, la thèse n'aborde pas complètement l'adaptation comportementale ou la réponse institutionnelle aux événements extrêmes et au changement climatique. Alors qu'elle se concentre sur les impacts de ces événements sur les communautés et les entreprises en général, elle n'aborde pas la manière dont les individus, les entreprises ou les gouvernements peuvent s'adapter et ainsi atténuer certains des effets négatifs. Il omet ainsi un facteur clé susceptible d'influencer la réponse aux chocs, en particulier dans les pays à revenu faible ou intermédiaire où les capacités d'adaptation peuvent varier considérablement.

Troisièmement, bien que la thèse introduise un nouveau jeu de données maillé permettant l'étude des événements multirisques, celui-ci demeure limité. Il ne contient ni rapports de dommages ni d'autres indicateurs potentiellement pertinents pour les décideurs publics. S'il est facile de le relier à des données agrégées au niveau national, aucun schéma n'est fourni pour expliciter cette procédure. De plus, avec une résolution spatiale de 10 km et une résolution temporelle mensuelle, le jeu de données manque de la précision nécessaire pour certains types d'analyses.

Quatrièmement, bien que l'accent géographique mis sur l'Afrique subsaharienne et l'Indonésie soit justifié, il limite la validité externe de la thèse. Les résultats sont influencés par les aléas dominants dans ces régions, mais d'autres zones du monde peuvent être exposées à des compositions d'événements extrêmes très différentes. Ainsi, cer-

taines conclusions pourraient ne pas être généralisables à l'échelle mondiale. Par exemple, les économies à revenu élevé pourraient présenter des dynamiques d'urbanisation différentes en réponse aux événements extrêmes, notamment en raison de systèmes d'assurance plus largement disponibles.

Enfin, la thèse ne traite pas directement de l'évolution de la fréquence et de l'intensité des événements étudiés. Ceux-ci sont considérés tels quels, bien que le changement climatique exerce une influence déterminante sur la distribution des événements extrêmes, non seulement sur leur fréquence mais aussi sur leur interconnexion.

Cette thèse ne fait qu'effleurer la surface des effets des événements extrêmes, en particulier lorsqu'ils sont interconnectés. Des questions fondamentales demeurent : Où les événements composés et cumulatifs surviennent-ils le plus fréquemment ? Quelles sont leurs conséquences socioéconomiques ? Sur quelles temporalités leurs effets persistent-ils, et par quels canaux se propagent-ils ? Par ailleurs, une définition plus rigoureuse et cohérente des événements composés est nécessaire afin d'aligner les approches théoriques et empiriques en modélisation économique.

Bien qu'il soit difficile de capturer pleinement la complexité des systèmes météorologiques et climatiques—en particulier en ce qui concerne les événements extrêmes et leurs impacts économiques—il est essentiel que les recommandations politiques fondées sur des données probantes reflètent la réalité aussi fidèlement que possible. Dans le cas contraire, nous risquons de concevoir des politiques inadéquates qui, dans le pire des cas, pourraient entraîner une mauvaise allocation des ressources nécessaires et exposer les populations et les acteurs économiques à des risques évitables.

Dans ce contexte, la collaboration interdisciplinaire est cruciale pour favoriser l'échange d'idées et de perspectives. Bien que cette thèse tente de contribuer à une meilleure compréhension des effets économiques des événements extrêmes, notamment dans les pays du Sud global, l'avenir offre des perspectives de recherche encore plus prometteuses. Comprendre comment les communautés et les sociétés réagissent à la menace croissante des événements extrêmes—qu'ils soient interconnectés ou isolés—n'est pas qu'un enjeu académique. C'est une condition centrale pour atteindre un développement équitable et durable au XXIeme siècle. Reste à savoir si ce sujet recevra l'attention qu'il mérite véritablement.

Bibliography

Abatzoglou, J. T., S. Z. Dobrowski, S. A. Parks, and K. C. Hegewisch (2018). "TerraClimate, a High-Resolution Global Dataset of Monthly Climate and Climatic Water Balance from 1958–2015". *Scientific Data* 5.1, p. 170191. DOI: 10.1038/sdata.2017.191.

Acemoglu, D., U. Akcigit, and W. Kerr (2016). "Networks and the Macroeconomy: An Empirical Exploration". *NBER Macroeconomics Annual* 30.1, pp. 273–335. DOI: 10.1086/685961.

Adger, W. N., A.-S. Crépin, C. Folke, D. Ospina, F. S. Chapin, K. Segerson, K. C. Seto, J. M. Anderies, S. Barrett, E. M. Bennett, G. Daily, T. Elmqvist, J. Fischer, N. Kautsky, S. A. Levin, J. F. Shogren, J. Van Den Bergh, B. Walker, and J. Wilen (2020). "Urbanization, Migration, and Adaptation to Climate Change". *One Earth* 3.4, pp. 396–399. DOI: 10.1016/j.oneear.2020.09.016.

AghaKouchak, A., F. Chiang, L. S. Huning, C. A. Love, I. Mallakpour, O. Mazdiyasni, H. Moftakhari, S. M. Papalexiou, E. Ragno, and M. Sadegh (2020). "Climate Extremes and Compound Hazards in a Warming World". *Annual Review of Earth and Planetary Sciences* 48.1, pp. 519–548. DOI: 10.1146/annurev-earth-071719-055228.

Ahrens, A., C. Hansen, and M. Schaffer (2024). *PDSLASSO: Stata Module for Post-Selection and Post-Regularization OLS or IV Estimation and Inference*. URL: https://EconPapers.repec.org/RePEc:boc:bocode:s458459. Pre-published.

AidData (2017). *WorldBank GeocodedResearchRelease Level1 v1.4.2 Geocoded Dataset*. URL: http://aiddata.org/research-datasets. (Accessed on 05/05/2024).

Albert, C., P. Bustos, and J. Ponticelli (2021). *The Effects of Climate Change on Labor and Capital Reallocation*. w28995. Cambridge, MA: National Bureau of Economic Research, w28995. DOI: 10.3386/w28995.

Amaral, P., J. Gaboardi, S. Rey, L. J. Wolf, W. Kang, P. Estrada, P. Stephens, eli knaap, D. C. Folch, D. Arribas-Bel, N. Malizia, C. Schmidt, D. Schult, J. Laura, J. Parry, T. Oshan, and X. Li (2025). *Pysal/Spreg: V1.8.3*. Version v1.8.3. Zenodo. DOI: 10.5281/ZENODO.15476723.

Anderson, W., C. Taylor, S. McDermid, E. Ilboudo-Nébié, R. Seager, W. Schlenker, F. Cottier, A. De Sherbinin, D. Mendeloff, and K. Markey (2021). "Violent Conflict Exacerbated Drought-Related Food Insecurity between 2009 and 2019 in Sub-Saharan Africa". *Nature Food* 2.8, pp. 603–615. DOI: 10.1038/s43016-021-00327-4.

Andrée, B. P. J., A. Chamorro, A. Kraay, P. Spencer, and D. Wang (2020). *Predicting Food Crises 2020, Dataset for Reproducing Working Paper Results*. WLD_2020_PFC_v01_M. World Bank Data Catalog. DOI: 10.48529/bwmq-rh91.

Anselin, L. (1988). *Spatial Econometrics: Methods and Models*. Vol. 4. Studies in Operational Regional Science. Dordrecht: Springer Netherlands. ISBN: 978-90-481-8311-1 978-94-015-7799-1. DOI: 10.1007/978-94-015-7799-1.

Antonioli, D., A. Marzucchi, and M. Modica (2022). "Resilience, Performance and Strategies in Firms' Reactions to the Direct and Indirect Effects of a Natural Disaster". *Networks and Spatial Economics*, pp. 1–25. DOI: 10.1007/s11067-021-09521-0.

Anuchitworawong, C. and K. Thampanishvong (2015). "Determinants of Foreign Direct Investment in Thailand: Does Natural Disaster Matter?" *International Journal of Disaster Risk Reduction* 14, pp. 312–321. DOI: 10.1016/j.ijdrr.2014.09.001.

Aragón, F. M., F. Oteiza, and J. P. Rud (2021). "Climate Change and Agriculture: Subsistence Farmers' Response to Extreme Heat". *American Economic Journal: Economic Policy* 13.1, pp. 1–35. DOI: 10.1257/pol.20190316.

Azevedo, J. P. (2011). *WBOPENDATA: Stata Module to Access World Bank Databases*. Version S457234. Boston College Department of Economics.

Badan Nasional Penanggulangan Bencana (2022). *Tabel Pencarian Data*. URL: https://gis.bnpb.go.id/. (Accessed on 01/05/2023).

Badan Pusat Statistik (2021). *Tabel Interregional Input-Output Indonesia Tahun 2016 Tahun Anggaran 2021*. 07100.2113. Badan Pusat Statistik. URL: https://www.bps.go.id/publication/2021/12/29/3ea49c0d856eceaba836792d/tabel-interregional-input-output-indonesia-tahun-2016-tahun-anggaran-2021.html. (Accessed on 02/02/2023).

Baker, A., B. Callaway, S. Cunningham, A. Goodman-Bacon, and P. H. C. Sant'Anna (2025). *Difference-in-Differences Designs: A Practitioner's Guide*. Version 1. DOI: 10.48550/ARXIV.2503.13323. Pre-published.

Batala, L. K., W. Yu, A. Khan, K. Regmi, and X. Wang (2021). "Natural Disasters' Influence on Industrial Growth, Foreign Direct Investment, and Export Performance in the South Asian Region of Belt and Road Initiative". *Natural Hazards* 108.2, pp. 1853–1876. DOI: 10.1007/s11069-021-04759-w.

Beck, H. E., T. R. McVicar, N. Vergopolan, A. Berg, N. J. Lutsko, A. Dufour, Z. Zeng, X. Jiang, A. I. J. M. Van Dijk, and D. G. Miralles (2023). "High-Resolution (1 Km) Köppen-Geiger

Maps for 1901–2099 Based on Constrained CMIP6 Projections". *Scientific Data* 10.1, p. 724. DOI: 10.1038/s41597-023-02549-6.

Beine, M. and L. Jeusette (2021). "A Meta-Analysis of the Literature on Climate Change and Migration". *Journal of Demographic Economics* 87.3, pp. 293–344. DOI: 10.1017/dem.2019.22.

Bellemare, M. F. and C. J. Wichman (2020). "Elasticities and the Inverse Hyperbolic Sine Transformation". *Oxford Bulletin of Economics and Statistics* 82.1, pp. 50–61. DOI: 10.1111/obes.12325.

Belloni, A., V. Chernozhukov, and C. Hansen (2014). "High-Dimensional Methods and Inference on Structural and Treatment Effects". *Journal of Economic Perspectives* 28.2, pp. 29–50. DOI: 10.1257/jep.28.2.29.

Berlemann, M. and M. F. Steinhardt (2017). "Climate Change, Natural Disasters, and Migration—a Survey of the Empirical Evidence". *CESifo Economic Studies* 63.4, pp. 353–385. DOI: 10.1093/cesifo/ifx019.

Bertoli, S., F. Docquier, H. Rapoport, and I. Ruyssen (2022). "Weather Shocks and Migration Intentions in Western Africa: Insights from a Multilevel Analysis". *Journal of Economic Geography* 22.2, pp. 289–323. DOI: 10.1093/jeg/lbab043.

Bettencourt, L. M. A. and N. Marchio (2023). *Million Neighborhoods Africa Database*. University of Chicago's Mansueto Institute for Urban Innovation. URL: https://www.millionneighborhoods.africa. (Accessed on 04/20/2025).

Blanc, E. and W. Schlenker (2017). "The Use of Panel Models in Assessments of Climate Impacts on Agriculture". *Review of Environmental Economics and Policy* 11.2, pp. 258–279. DOI: 10.1093/reep/rex016.

Blanc, E. and E. Strobl (2016). "Assessing the Impact of Typhoons on Rice Production in the Philippines". *Journal of Applied Meteorology and Climatology* 55.4, pp. 993–1007. DOI: 10.1175/JAMC-D-15-0214.1.

Blocher, J. M., R. Hoffmann, and H. Weisz (2024). "The Effects of Environmental and Non-Environmental Shocks on Livelihoods and Migration in Tanzania". *Population and Environment* 46.1, p. 7. DOI: 10.1007/s11111-024-00449-4.

Boehm, C. E., A. Flaaen, and N. Pandalai-Nayar (2019). "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake". *The Review of Economics and Statistics* 101.1, pp. 60–75. DOI: 10.1162/rest_a_00750.

Bohra-Mishra, P., M. Oppenheimer, and S. Hsiang (2014). "Nonlinear Permanent Migration Response to Climatic Variations but Minimal Response to Disasters". *Proceedings of the National Academy of Sciences* 111.27, pp. 9780–9785. DOI: 10.1073/pnas.1317166111.

Borusyak, K. (2021). *EVENT_PLOT: Stata Module to Plot the Staggered-adoption Diff-in-Diff (" Event Study") Estimates*. Pre-published.

Borusyak, K., X. Jaravel, and J. Spiess (2024). "Revisiting Event Study Designs: Robust and Efficient Estimation". *The Review of Economic Studies* 91.6, pp. 3253–3285. DOI: 10.1093/restud/rdae007.

Bosker, M., S. Brakman, H. Garretsen, and M. Schramm (2007). "Looking for Multiple Equilibria When Geography Matters: German City Growth and the WWII Shock". *Journal of Urban Economics* 61.1, pp. 152–169. DOI: 10.1016/j.jue.2006.07.001.

Botzen, W. J. W., O. Deschenes, and M. Sanders (2019). "The Economic Impacts of Natural Disasters: A Review of Models and Empirical Studies". *Review of Environmental Economics and Policy* 13.2, pp. 167–188. DOI: 10.1093/reep/rez004.

Brakenbridge, G. and Dartmouth Flood Observatory (2023). *Global Active Archive of Large Flood Events*. URL: https://floodobservatory.colorado.edu/Archives/index.html. (Accessed on 03/27/2023).

Breitung, J., S. Kripfganz, and K. Hayakawa (2022). "Bias-Corrected Method of Moments Estimators for Dynamic Panel Data Models". *Econometrics and Statistics* 24, pp. 116–132. DOI: 10.1016/j.ecosta.2021.07.001.

Brucal, A. and S. Mathews (2021). *Market Entry, Survival, and Exit of Firms in the Aftermath of Natural Hazard-related Disasters: A Case Study of Indonesian Manufacturing Plants*. Pre-published.

Bukvic, A. and S. Barnett (2023). "Drivers of Flood-Induced Relocation among Coastal Urban Residents: Insight from the US East Coast". *Journal of Environmental Management* 325, p. 116429. DOI: 10.1016/j.jenvman.2022.116429.

Burke, M., A. Driscoll, S. Heft-Neal, J. Xue, J. Burney, and M. Wara (2021). "The Changing Risk and Burden of Wildfire in the United States". *Proceedings of the National Academy of Sciences* 118.2, e2011048118. DOI: 10.1073/pnas.2011048118.

Burke, M., S. Hsiang, and E. Miguel (2015). "Global Non-Linear Effect of Temperature on Economic Production". *Nature* 527.7577, pp. 235–239. DOI: 10.1038/nature15725.

Burke, M., M. Zahid, M. C. Martins, C. Callahan, R. Lee, T. Avirmed, S. Heft-Neal, M. Kiang, S. Hsiang, and D. Lobell (2024). *Are We Adapting to Climate Change?* w32985. Cambridge, MA: National Bureau of Economic Research, w32985. DOI: 10.3386/w32985.

Callaway, B. and P. H. C. Sant'Anna (2021). "Difference-in-Differences with Multiple Time Periods". *Journal of Econometrics* 225.2, pp. 200–230. DOI: 10.1016/j.jeconom.2020.12.001.

Camacho, R., J. Aryal, and A. Rajabifard (2025). "How Do Disasters Disrupt the Spatial Growth of Informal Settlements? A Multi-Temporal Remote Sensing Approach – The Case Study of Mocoa, Colombia". *Habitat International* 156, p. 103272. DOI: 10.1016/j.habitatint.2024.103272.

Cameron, G. J., H.-A. H. Dang, M. Dinc, J. Foster, and M. M. Lokshin (2021). "Measuring the Statistical Capacity of Nations*". *Oxford Bulletin of Economics and Statistics* 83.4, pp. 870–896. DOI: 10.1111/obes.12421.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021). "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake". *Quarterly Journal of Economics* 136.2, pp. 1255–1321. DOI: 10.1093/qje/qjaa044.

Castells-Quintana, D., M. Krause, and T. K. J. McDermott (2021). "The Urbanising Force of Global Warming: The Role of Climate Change in the Spatial Distribution of Population". *Journal of Economic Geography* 21.4, pp. 531–556. DOI: 10.1093/jeg/lbaa030.

Castells-Quintana, D., M. D. P. Lopez-Uribe, and T. K. McDermott (2022). "Population Displacement and Urban Conflict: Global Evidence from More than 3300 Flood Events". *Journal of Development Economics* 158, p. 102922. DOI: 10.1016/j.jdeveco.2022.102922.

Cattaneo, C., M. Beine, C. J. Fröhlich, D. Kniveton, I. Martinez-Zarzoso, M. Mastrorillo, K. Millock, E. Piguet, and B. Schraven (2019). "Human Migration in the Era of Climate Change". *Review of Environmental Economics and Policy* 13.2, pp. 189–206. DOI: 10.1093/reep/rez008.

Cattaneo, C. and G. Peri (2016). "The Migration Response to Increasing Temperatures". *Journal of Development Economics* 122, pp. 127–146. DOI: 10.1016/j.jdeveco.2016.05.004.

Cevik, S. and J. Jalles (2024). "Eye of the Storm: The Impact of Climate Shocks on Inflation and Growth". *Review of Economics* 75.2, pp. 109–138. DOI: 10.1515/roe-2024-0005.

Chen, J. J., V. Mueller, Y. Jia, and S. K.-H. Tseng (2017). "Validating Migration Responses to Flooding Using Satellite and Vital Registration Data". *American Economic Review* 107.5, pp. 441–445. DOI: 10.1257/aer.p20171052.

Chen, Z., B. Yu, C. Yang, Y. Zhou, S. Yao, X. Qian, C. Wang, B. Wu, and J. Wu (2021). "An Extended Time Series (2000–2018) of Global NPP-VIIRS-like Nighttime Light Data from a Cross-Sensor Calibration". *Earth System Science Data* 13.3, pp. 889–906. DOI: 10.5194/essd-13-889-2021.

Chi, G., H. Fang, S. Chatterjee, and J. E. Blumenstock (2022). "Microestimates of Wealth for All Low- and Middle-Income Countries". *Proceedings of the National Academy of Sciences* 119.3, e2113658119. DOI: 10.1073/pnas.2113658119.

Chlouba, V., M. Mukim, and E. Zaveri (2023). *After Big Droughts Come Big Cities: Does Drought Drive Urbanization?* 10408. World Bank, Washington, DC. URL: https://openknowledge.worldbank.org/entities/publication/53e73155-a428-4d0f-a22e-4242d3ed9f20. (Accessed on 08/26/2024).

Claassen, J. N., P. J. Ward, J. Daniell, E. E. Koks, T. Tiggeloven, and M. C. De Ruiter (2023a). "A New Method to Compile Global Multi-Hazard Event Sets". *Scientific Reports* 13.1, p. 13808. DOI: 10.1038/s41598-023-40400-5.

— (2023b). *MYRIAD – Hazard Event Sets (MYRIAD-HES)*. Zenodo. DOI: 10.5281/ZENODO.8269679. (Accessed on 03/25/2024).

Clark, J., C. E. Konrad, and A. Grundstein (2024). "The Development and Accuracy Assessment of Wet Bulb Globe Temperature Forecasts". *Weather and Forecasting* 39.2, pp. 403–419. DOI: 10.1175/WAF-D-23-0076.1.

Climate Hazards Center (UCSB) (2023). *CHIRPS: Rainfall Estimates from Rain Gauge and Satellite Observations*. URL: https://chc.ucsb.edu/data/chirps/. (Accessed on 06/29/2023).

Cole, M. A., R. J. R. Elliott, T. Okubo, and E. Strobl (2019). "Natural Disasters and Spatial Heterogeneity in Damages: The Birth, Life and Death of Manufacturing Plants". *Journal of Economic Geography* 19.2, pp. 373–408. DOI: 10.1093/jeg/lbx037.

Colella, F., R. Lalive, S. O. Sakalli, and M. Thoenig (2019). "Inference with Arbitrary Clustering". *SSRN Electronic Journal*. DOI: 10.2139/ssrn.3449578.

Colmer, J. (2021). "Temperature, Labor Reallocation, and Industrial Production: Evidence from India". *American Economic Journal: Applied Economics* 13.4, pp. 101–124. DOI: 10.1257/app.20190249.

Conigliani, C., V. Costantini, and G. Finardi (2022). "Climate-Related Natural Disasters and Forced Migration: A Spatial Regression Analysis". *Spatial Economic Analysis* 17.3, pp. 416–439. DOI: 10.1080/17421772.2021.1995620.

Copernicus Climate Change Service (2019). *Fire Burned Area from 2001 to Present Derived from Satellite Observations*. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). DOI: 10.24381/CDS.F333CF85. (Accessed on 03/27/2023).

— (2022). *ERA5-Land Monthly Averaged Data from 1950 to Present*. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). DOI: 10.24381/cds.68d2bb30. (Accessed on 06/29/2023).

— (2024a). *ERA5-Land Post-Processed Daily Statistics from 1950 to Present*. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). DOI: 10.24381/cds.e9c9c792. (Accessed on 10/18/2024).

— (2024b). *Daily Statistics Calculated from ERA5 Data*. URL: https://cds.climate.copernicus.eu/cdsapp#!/software/app-c3s-daily-era5-statistics?tab=overview. (Accessed on 07/11/2024).

Crameri, F., G. E. Shephard, and P. J. Heron (2020). "The Misuse of Colour in Science Communication". *Nature Communications* 11.1, p. 5444. DOI: 10.1038/s41467-020-19160-7.

CRED and UC Louvain (2024). *EM-DAT: International Disaster Database*. URL: www.emdat.be. (Accessed on 09/13/2024).

Damette, O. and S. Goutte (2023). "Beyond Climate and Conflict Relationships: New Evidence from a Copula-based Analysis on an Historical Perspective". *Journal of Comparative Economics* 51.1, pp. 295–323. DOI: 10.1016/j.jce.2022.09.005.

Data for Good at Meta (2023). *Facebook Commuting Zones*. URL: https://dataforgood.facebook.com/dfg/tools/commuting-zones. (Accessed on 04/28/2023).

Davis, D. R. and D. E. Weinstein (2002). "Bones, Bombs, and Break Points: The Geography of Economic Activity". *American Economic Review* 92.5, pp. 1269–1289. DOI: 10.1257/000282802762024502.

De, P. K. and D. Thamarapani (2022). "Impacts of Negative Shocks on Wellbeing and Aspirations – Evidence from an Earthquake". *World Development* 154, p. 105876. DOI: 10.1016/j.worlddev.2022.105876.

De Chaisemartin, C. and X. d'Haultfoeuille (2020). "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects". *American Economic Review* 110.9, pp. 2964–2996. DOI: 10.1257/aer.20181169.

— (2023). "Two-Way Fixed Effects and Differences-in-Differences Estimators with Several Treatments". *Journal of Econometrics* 236.2, p. 105480. DOI: 10.1016/j.jeconom.2023.105480.

— (2024). "Difference-in-Differences Estimators of Intertemporal Treatment Effects". *Review of Economics and Statistics*, pp. 1–45. DOI: 10.1162/rest_a_01414.

De Mel, S., D. McKenzie, and C. Woodruff (2012). "Enterprise Recovery Following Natural Disasters". *The Economic Journal* 122.559, pp. 64–91. DOI: 10.1111/j.1468-0297.2011.02475.x.

Dell, M., B. F. Jones, and B. A. Olken (2012). "Temperature Shocks and Economic Growth: Evidence from the Last Half Century". *American Economic Journal: Macroeconomics* 4.3, pp. 66–95. DOI: 10.1257/mac.4.3.66.

Dell, M., B. F. Jones, and B. A. Olken (2014). "What Do We Learn from the Weather? The New Climate-Economy Literature". *Journal of Economic Literature* 52.3, pp. 740–798. DOI: 10.1257/jel.52.3.740.

Deschênes, O. and M. Greenstone (2007). "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather". *American Economic Review* 97.1, pp. 354–385. DOI: 10.1257/aer.97.1.354.

Di Falco, S., A. B. Kis, M. Viarengo, and U. Das (2024). "Leaving Home: Cumulative Climate Shocks and Migration in Sub-Saharan Africa". *Environmental and Resource Economics* 87.1, pp. 321–345. DOI: 10.1007/s10640-023-00826-x.

Doan, M. K., R. Hill, S. Hallegatte, P. Corral, B. Brunckhorst, M. Nguyen, S. Freije-Rodriguez, and E. Naikal (2023). *Counting People Exposed to, Vulnerable to, or at High Risk From Climate Shocks: A Methodology*. World Bank, Washington, DC. DOI: 10.1596/1813-9450-10619.

Dobes, L., F. Jotzo, and D. I. Stern (2014). "The Economics of Global Climate Change: A Historical Literature Review". *Review of Economics* 65.3, pp. 281–320. DOI: 10.1515/roe-2014-0305.

Dodman, D., B. Hayward, M. Pelling, V. Castan Broto, W. Chow, E. Chu, R. Dawson, L. Khirfan, T. McPhearson, A. Prakash, Y. Zheng, and G. Ziervogel (2022). "Cities, Settlements and Key Infrastructure". *Climate Change 2022: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by H.-O. Pörtner, D. Roberts, M. Tignor, E. Poloczanska, K. Mintenbeck, A. Alegria, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, and B. Rama. Cambridge, UK and New York, USA: Cambridge University Press, pp. 907–1040. URL: https://doi.org/10.1017/9781009325844.008.

Done, J. (2022). *Willis Research Network Global Tropical Cyclone Footprint Dataset*. figshare. DOI: 10.6084/M9.FIGSHARE.19178426.V1. (Accessed on 07/11/2024).

Done, J. M., M. Ge, G. J. Holland, I. Dima-West, S. Phibbs, G. R. Saville, and Y. Wang (2020). "Modelling Global Tropical Cyclone Wind Footprints". *Natural Hazards and Earth System Sciences* 20.2, pp. 567–580. DOI: 10.5194/nhess-20-567-2020.

Dormady, N., A. Roa-Henriquez, and A. Rose (2019). "Economic Resilience of the Firm: A Production Theory Approach". *International Journal of Production Economics* 208, pp. 446–460. DOI: 10.1016/j.ijpe.2018.07.017.

Doytch, N. (2019). "Upgrading Destruction?" *International Journal of Climate Change Strategies and Management* 12.2, pp. 182–200. DOI: 10.1108/IJCCSM-07-2019-0044.

Drabo, A. (2021). "How Do Climate Shocks Affect the Impact of FDI, ODA and Remittances on Economic Growth?" *IMF Working Papers* 2021.193, p. 1. DOI: 10.5089/9781513585635.001.

Eberenz, S., D. Stocker, T. Röösli, and D. N. Bresch (2020). "Asset Exposure Data for Global Physical Risk Assessment". *Earth System Science Data* 12.2, pp. 817–833. DOI: 10.5194/essd-12-817-2020.

Eberle, U. J., D. Rohner, and M. Thoenig (2025). "Heat and Hate: Climate Security and Farmer-Herder Conflicts in Africa". *Review of Economics and Statistics*, pp. 1–47. DOI: 10.1162/rest_a_01546.

Elliott, R. J. R., Y. Liu, E. Strobl, and M. Tong (2019). "Estimating the Direct and Indirect Impact of Typhoons on Plant Performance: Evidence from Chinese Manufacturers". *Journal of Environmental Economics and Management* 98, p. 102252. DOI: 10.1016/J.JEEM.2019.102252.

Elvidge, C. D., K. E. Baugh, E. A. Kihn, H. W. Kroehl, E. R. Davis, and C. W. Davis (1997). "Relation between Satellite Observed Visible-near Infrared Emissions, Population, Economic Activity and Electric Power Consumption". *International Journal of Remote Sensing* 18.6, pp. 1373–1379. DOI: 10.1080/014311697218485.

Emrich, C. T., Y. Zhou, S. K. Aksha, and H. E. Longenecker (2022). "Creating a Nationwide Composite Hazard Index Using Empirically Based Threat Assessment Approaches Applied to Open Geospatial Data". *Sustainability* 14.5. DOI: 10.3390/su14052685.

Esteban, M., V. Tsimopoulou, T. Mikami, N. Yun, A. Suppasri, and T. Shibayama (2013). "Recent Tsunamis Events and Preparedness: Development of Tsunami Awareness in Indonesia, Chile and Japan". *International Journal of Disaster Risk Reduction* 5, pp. 84–97. DOI: 10.1016/j.ijdrr.2013.07.002.

European Commission - JRC and International Energy Agency (IEA) (2023). *EDGAR (Emissions Database for Global Atmospheric Research) Community GHG Database - Version 8.0*. European Commission, JRC (Datasets). URL: https://edgar.jrc.ec.europa.eu/report_2023. (Accessed on 04/15/2024).

European Severe Storms Laboratory (2014). *ESWD Event Reporting Criteria*. European Severe Weather Database. URL: https://www.eswd.eu/docs/ESWD_criteria_en.pdf. (Accessed on 05/08/2024).

Felbermayr, G. and J. Gröschl (2014). "Naturally Negative: The Growth Effects of Natural Disasters". *Journal of Development Economics* 111, pp. 92–106. DOI: 10.1016/j.jdeveco.2014.07.004.

— (2024). *Ifo GAME – the Geological and Meteorological Events Database*. URL: https://www.ifo.de/en/ebdc-dataset/ifo-game-geological-and-meteorological-events-database. (Accessed on 10/18/2024).

Felbermayr, G., J. Gröschl, M. Sanders, V. Schippers, and T. Steinwachs (2022). "The Economic Impact of Weather Anomalies". *World Development* 151. DOI: 10.1016/j.worlddev.2021.105745.

Feng, B., Y. Zhang, and R. Bourke (2021). "Urbanization Impacts on Flood Risks Based on Urban Growth Data and Coupled Flood Models". *Natural Hazards* 106.1, pp. 613–627. DOI: 10.1007/s11069-020-04480-0.

Ferreira, S. (2024). "Extreme Weather Events and Climate Change: Economic Impacts and Adaptation Policies". *Annual Review of Resource Economics* 16.1, pp. 207–231. DOI: 10.1146/annurev-resource-101623-095314.

Fezzi, C. and I. Bateman (2015). "The Impact of Climate Change on Agriculture: Nonlinear Effects and Aggregation Bias in Ricardian Models of Farmland Values". *Journal of the Association of Environmental and Resource Economists* 2.1, pp. 57–92. DOI: 10.1086/680257.

Food and Agriculture Organization (FAO) of the United Nations (2023). *Locust Hub*. URL: https://locust-hub-hqfao.hub.arcgis.com/. (Accessed on 05/26/2023).

Founda, D. and M. Santamouris (2017). "Synergies between Urban Heat Island and Heat Waves in Athens (Greece), during an Extremely Hot Summer (2012)". *Scientific Reports* 7.1, p. 10973. DOI: 10.1038/s41598-017-11407-6.

Friedt, F. L. and A. Toner-Rodgers (2022). "Natural Disasters, Intra-national FDI Spillovers, and Economic Divergence: Evidence from India". *Journal of Development Economics* 157. DOI: 10.1016/j.jdeveco.2022.102872.

Froude, M. J. and D. N. Petley (2018). "Global Fatal Landslide Occurrence from 2004 to 2016". *Natural Hazards and Earth System Sciences* 18.8, pp. 2161–2181. DOI: 10.5194/nhess-18-2161-2018.

Funk, C., P. Peterson, M. Landsfeld, D. Pedreros, J. Verdin, J. Rowland, B. Romero, G. Husak, J. Michaelsen, and A. Verdin (2014). "A Quasi-Global Precipitation Time Series for Drought Monitoring". *U.S. Geological Survey Data Series* 832, p. 4. DOI: 10.3133/ds832.

Funk, C., P. Peterson, M. Landsfeld, D. Pedreros, J. Verdin, S. Shukla, G. Husak, J. Rowland, L. Harrison, A. Hoell, and J. Michaelsen (2015). "The Climate Hazards Infrared Precipitation with Stations—a New Environmental Record for Monitoring Extremes". *Scientific Data* 2.1, p. 150066. DOI: 10.1038/sdata.2015.66.

Gahtan, J., K. Knapp, C. Schreck, H. Diamond, J. Kossin, and M. Kruk (2018). *International Best Track Archive for Climate Stewardship (IBTrACS) Project, Version 4*. NOAA National Centers for Environmental Information. DOI: 10.25921/82ty-9e16.

Galbusera, L. and G. Giannopoulos (2018). "On Input-output Economic Models in Disaster Impact Assessment". *International Journal of Disaster Risk Reduction* 30, pp. 186–198. DOI: 10.1016/j.ijdrr.2018.04.030.

Gandhi, S., M. Kahn, R. Kochhar, S. Lall, and V. Tandel (2022). *Adapting to Flood Risk: Evidence from a Panel of Global Cities*. w30137. Cambridge, MA: National Bureau of Economic Research, w30137. DOI: 10.3386/w30137.

Gebrechorkos, S., J. Peng, E. Dyer, D. G. Miralles, S. M. Vicente-Serrano, C. Funk, H. Beck, D. Asfaw, M. Singer, and S. Dadson (2023). *Hydro-JULES: Global High-Resolution Drought Datasets from 1981-2022*. application/xml. NERC EDS Centre for Environmental Data Analysis. DOI: 10.5285/AC43DA11867243A1BB414E1637802DEC. (Accessed on 10/01/2024).

GeoNames (n.d.). *GeoNames Database*. URL: https://www.geonames.org/.

German Aerospace Center (DLR) (2024). *World Settlement Footprint (WSF) Evolution - Landsat-5/-7 - Global*. URL: https://download.geoservice.dlr.de/WSF_EVO/#details. (Accessed on 08/15/2024).

Ghanbari, M., M. Arabi, M. Georgescu, and A. M. Broadbent (2023). "The Role of Climate Change and Urban Development on Compound Dry-Hot Extremes across US Cities". *Nature Communications* 14.1, p. 3509. DOI: 10.1038/s41467-023-39205-x.

Gibson, J., Y. Jiang, X. Zhang, and G. Boe-Gibson (2024). "Are Disaster Impact Estimates Distorted by Errors in Popular Night-Time Lights Data?" *Economics of Disasters and Climate Change* 8.3, pp. 391–416. DOI: 10.1007/s41885-024-00152-6.

Gibson, J., B. Kim, and C. Li (2024). *Luminosity and Local Economic Growth*. URL: https://EconPapers.repec.org/RePEc:wai:econwp:24/08. (Accessed on 12/13/2024). Pre-published.

Gignoux, J. and M. Menéndez (2016). "Benefit in the Wake of Disaster: Long-run Effects of Earthquakes on Welfare in Rural Indonesia". *Journal of Development Economics* 118, pp. 26–44. DOI: 10.1016/j.jdeveco.2015.08.004.

Giuseppe, F. D., C. Vitolo, B. Krzeminski, C. Barnard, and J. S. Miguel (2019). *Fire Weather Index - ERA-Interim*. Version 3.0. Zenodo. DOI: 10.5281/zenodo.3251000. (Accessed on 06/29/2023).

Global Volcanism Program and E. Venzke (2023). *Volcanoes of the World, v.5.1.0*. Global Volcanism Program. DOI: 10.5479/si.GVP.VOTW5-2023.5.1.

Gómez, D., E. F. García, and E. Aristizábal (2023). "Spatial and Temporal Landslide Distributions Using Global and Open Landslide Databases". *Natural Hazards* 117.1, pp. 25–55. DOI: 10.1007/s11069-023-05848-8.

Goodman, S., A. BenYishay, Z. Lv, and D. Runfola (2019). "GeoQuery: Integrating HPC Systems and Public Web-Based Geospatial Data Tools". *Computers & Geosciences* 122 (June 2018), pp. 103–112. DOI: 10.1016/j.cageo.2018.10.009.

Goodman-Bacon, A. (2021). "Difference-in-Differences with Variation in Treatment Timing". *Journal of Econometrics* 225.2, pp. 254–277. DOI: 10.1016/j.jeconom.2021.03.014.

Gröger, A. and Y. Zylberberg (2016). "Internal Labor Migration as a Shock Coping Strategy: Evidence from a Typhoon". *American Economic Journal: Applied Economics* 8.2, pp. 123–153. DOI: 10.1257/app.20140362.

Gu, G. W. and G. Hale (2023). "Climate Risks and FDI". *Journal of International Economics*, p. 103731. DOI: 10.1016/j.jinteco.2023.103731.

Güneralp, B., İ. Güneralp, and Y. Liu (2015). "Changing Global Patterns of Urban Exposure to Flood and Drought Hazards". *Global Environmental Change* 31, pp. 217–225. DOI: 10.1016/j.gloenvcha.2015.01.002.

Hallegatte, S. (2008). "An Adaptive Regional Input-Output Model and Its Application to the Assessment of the Economic Cost of Katrina". *Risk Analysis* 28.3, pp. 779–799. DOI: 10.1111/j.1539-6924.2008.01046.x.

Hallegatte, S. (2014). *Economic Resilience: Definition and Measurement*. World Bank, Washington, DC. DOI: 10.1596/1813-9450-6852.

— (2019). "Disasters' Impacts on Supply Chains". *Nature Sustainability* 2.9, pp. 791–792. DOI: 10.1038/s41893-019-0380-5.

Hallegatte, S. and P. Dumas (2009). "Can Natural Disasters Have Positive Consequences? Investigating the Role of Embodied Technical Change". *Ecological Economics* 68.3, pp. 777–786. DOI: 10.1016/j.ecolecon.2008.06.011.

Hallegatte, S., A. Vogt-Schilb, M. Bangalore, and J. Rozenberg (2017). *Unbreakable: Building the Resilience of the Poor in the Face of Natural Disasters*. Climate Change and Development. World Bank, Washington, DC. URL: http://documents.worldbank.org/curated/en/512241480487839624.

Hallegatte, S., A. Vogt-Schilb, J. Rozenberg, M. Bangalore, and C. Beaudet (2020). "From Poverty to Disaster and Back: A Review of the Literature". *Economics of Disasters and Climate Change* 4.1, pp. 223–247. DOI: 10.1007/s41885-020-00060-5.

Hamza, M., K. Eriksson, and R. Staupe-Delgado (2021). "Locating Potential Sources of Capacity and Vulnerability in Geographically Remote Areas: Reflections Based on Three Case Studies". *International Journal of Disaster Risk Reduction* 63, p. 102433. DOI: 10.1016/j.ijdrr.2021.102433.

HarvestChoice; International Food Policy Research Institute (IFPRI) (2016). *Travel Time to Markets in Africa South of the Sahara*. Harvard Dataverse. DOI: 10.7910/DVN/YKDWJD. (Accessed on 11/07/2023).

Heinen, A., J. Khadan, and E. Strobl (2019). "The Price Impact of Extreme Weather in Developing Countries". *The Economic Journal* 129.619, pp. 1327–1342. DOI: 10.1111/ecoj.12581.

Helbling, M. and D. Meierrieks (2023). "Global Warming and Urbanization". *Journal of Population Economics* 36.3, pp. 1187–1223. DOI: 10.1007/s00148-022-00924-y.

Henderson, V. and A. Kuncoro (2011). "Corruption and Local Democratization in Indonesia: The Role of Islamic Parties". *Journal of Development Economics* 94.2, pp. 164–180. DOI: 10.1016/j.jdeveco.2010.01.007.

Henderson, V., A. Storeygard, and U. Deichmann (2017). "Has Climate Change Driven Urbanization in Africa?" *Journal of Development Economics* 124, pp. 60–82. DOI: 10.1016/j.jdeveco.2016.09.001.

Henderson, V., A. Storeygard, and D. N. Weil (2011). "A Bright Idea for Measuring Economic Growth". *American Economic Review* 101.3, pp. 194–199. DOI: 10.1257/aer.101.3.194.

Hersbach, H., B. Bell, B. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, D. Schepers, A. Simmons, C. Soci, D. Dee, and J. Thépaut (2023). *ERA5 Hourly Data on Single Levels from 1940 to Present*. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). DOI: 10.24381/CDS.ADBB2D47. (Accessed on 10/18/2024).

Hoffmann, R., A. Dimitrova, R. Muttarak, J. Crespo Cuaresma, and J. Peisker (2020). "A Meta-Analysis of Country-Level Studies on Environmental Change and Migration". *Nature Climate Change* 10.10, pp. 904–912. DOI: 10.1038/s41558-020-0898-6.

Hoffmann, R., B. Šedová, and K. Vinke (2021). "Improving the Evidence Base: A Methodological Review of the Quantitative Climate Migration Literature". *Global Environmental Change* 71, p. 102367. DOI: 10.1016/j.gloenvcha.2021.102367.

Hogan, D. and W. Schlenker (2024). "Empirical Approaches to Climate Change Impact Quantification". *Handbook of the Economics of Climate Change*. Vol. 1. Elsevier, pp. 53–111. ISBN: 978-0-443-31324-0. DOI: 10.1016/bs.hesecc.2024.10.006.

Hornbeck, R. and D. Keniston (2017). "Creative Destruction: Barriers to Urban Growth and the Great Boston Fire of 1872". *American Economic Review* 107.6, pp. 1365–1398. DOI: 10.1257/aer.20141707.

Hsiang, S. (2016). "Climate Econometrics". *Annual Review of Resource Economics* 8.1, pp. 43–75. DOI: 10.1146/annurev-resource-100815-095343.

Hsiang, S. and A. Jina (2014). *The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones*. w20352. Cambridge, MA: National Bureau of Economic Research, w20352. DOI: 10.3386/w20352.

Hu, Y. and J. Yao (2022). "Illuminating Economic Growth". *Journal of Econometrics* 228.2, pp. 359–378. DOI: 10.1016/j.jeconom.2021.05.007.

Huang, R., A. Malik, M. Lenzen, Y. Jin, Y. Wang, F. Faturay, and Z. Zhu (2022). "Supply-Chain Impacts of Sichuan Earthquake: A Case Study Using Disaster Input–Output Analysis". *Natural Hazards* 110.3, pp. 2227–2248. DOI: 10.1007/s11069-021-05034-8.

Huang, W. T. K., P. Masselot, E. Bou-Zeid, S. Fatichi, A. Paschalis, T. Sun, A. Gasparrini, and G. Manoli (2023). "Economic Valuation of Temperature-Related Mortality Attributed to Urban Heat Islands in European Cities". *Nature Communications* 14.1, p. 7438. DOI: 10.1038/s41467-023-43135-z.

Imai, K. and I. S. Kim (2019). "When Should We Use Unit Fixed Effects Regression Models for Causal Inference with Longitudinal Data?" *American Journal of Political Science* 63.2, pp. 467–490. DOI: 10.1111/ajps.12417.

Imai, K., I. S. Kim, and E. Wang (2023). "Matching Methods for Causal Inference with Time-Series Cross-Sectional Data". *American Journal of Political Science* 67.3, pp. 587–605. DOI: 10.1111/ajps.12685.

Indonesian Investment Co-ordinating Board (BKPM) (2021). *Panduan Penggunaan*. URL: https://nswi.bkpm.go.id/. (Accessed on 03/27/2021).

Inness, A., M. Ades, A. Agustí-Panareda, J. Barré, A. Benedictow, A.-M. Blechschmidt, J. J. Dominguez, R. Engelen, H. Eskes, J. Flemming, V. Huijnen, L. Jones, Z. Kipling, S. Massart, M. Parrington, V.-H. Peuch, M. Razinger, S. Remy, M. Schulz, and M. Suttie (2019). "The CAMS Reanalysis of Atmospheric Composition". *Atmospheric Chemistry and Physics* 19.6, pp. 3515–3556. DOI: 10.5194/acp-19-3515-2019.

Inoue, H. and Y. Todo (2019). "Firm-Level Propagation of Shocks through Supply-Chain Networks". *Nature Sustainability* 2.9, pp. 841–847. DOI: 10.1038/s41893-019-0351-x.

International Centre for Geohazards and H. Smebye (2018). *Landslides - Frequency (Triggered by Precipitations)*. URL: https://web.archive.org/web/20220308150423/https://preview.grid.unep.ch/index.php?preview=data&events=landslides&evcat=2&lang=eng. (Accessed on 01/03/2021).

International Centre for Geohazards / NGI and F. Lovholt (2009). *Tsunami - Frequency*. URL: https://web.archive.org/web/20220308130428/https://preview.grid.unep.ch/index.php. (Accessed on 01/03/2021).

International Centre for Geohazards / NGI and H. Smebye (2018). *Landslides - Frequency (Triggered by Earthquakes)*. URL: https://web.archive.org/web/20220121173835/https://preview.grid.unep.ch/index.php?preview=data&events=landslides&evcat=1&lang=eng. (Accessed on 01/03/2021).

International Food Policy Research Institute (2020). *Spatially-Disaggregated Crop Production Statistics Data in Africa South of the Sahara for 2017*. Harvard Dataverse. DOI: 10.7910/DVN/FSSKBW. (Accessed on 06/29/2023).

IPCC (2012). *Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation*. Ed. by C. Field, V. Barros, T. Stocker, D. Qin, D. Dokken, K. Ebi, M. Mastrandrea, K. Mach, G.-K. Plattner, S. Allen, M. Tignor, and P. Midgley. A Special Report of Working Groups I and II of the Intergovernmental Panel on Climate Chang. Cambridge University Press. 582pp. URL: https://www.ipcc.ch/report/managing-the-risks-of-extreme-events-and-disasters-to-advance-climate-change-adaptation/.

— (2021). *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by V. Masson-Delmotte, P. Zhai, A. Pirani, S. Connors, C. Pean, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. Matthews, T. Maycock, T. Waterfield, O. Yelekci, R. Yu, and B. Zhou. 1st ed. Cambridge, United Kingdom and New York, NY: Cambridge University Press. ISBN: 978-1-00-915789-6. DOI: 10.1017/9781009157896.

Javorcik, B. S., Ç. Özden, M. Spatareanu, and C. Neagu (2011). "Migrant Networks and Foreign Direct Investment". *Journal of Development Economics* 94.2, pp. 231–241. DOI: 10.1016/j.jdeveco.2010.01.012.

Jiménez Martínez, M., M. Jiménez Martínez, and R. Romero-Jarén (2020). "How Resilient Is the Labour Market against Natural Disaster? Evaluating the Effects from the 2010 Earthquake in Chile". *Natural Hazards* 104.2, pp. 1481–1533. DOI: 10.1007/s11069-020-04229-9.

Joint Research Center and Copernicus Emergency Management Service (2019). *River Discharge Historical Data from the Global Flood Awareness System*. Early Warning Data Store (EWDS). DOI: 10.24381/CDS.A4FDD6B9. (Accessed on 01/24/2025).

Kaczan, D. J. and J. Orgill-Meyer (2020). "The Impact of Climate Change on Migration: A Synthesis of Recent Empirical Insights". *Climatic Change* 158.3-4, pp. 281–300. DOI: 10.1007/s10584-019-02560-0.

Kaltenböck, R., G. Diendorfer, and N. Dotzek (2009). "Evaluation of Thunderstorm Indices from ECMWF Analyses, Lightning Data and Severe Storm Reports". *Atmospheric Research* 93.1-3, pp. 381–396. DOI: 10.1016/j.atmosres.2008.11.005.

Kato, H. and T. Okubo (2022). "The Resilience of FDI to Natural Disasters Through Industrial Linkages". *Environmental and Resource Economics* 82.1, pp. 177–225. DOI: 10.1007/s10640-022-00666-1.

Katoka, B. (2021). "Do Natural Disasters Reduce Foreign Direct Investment in Sub-Saharan Africa?" *Economic Effects of Natural Disasters*, pp. 529–546. DOI: 10.1016/B978-0-12-817465-4.00031-5.

Kiely, L., D. V. Spracklen, S. R. Arnold, E. Papargyropoulou, L. Conibear, C. Wiedinmyer, C. Knote, and H. A. Adrianto (2021). "Assessing Costs of Indonesian Fires and the Benefits of Restoring Peatland". *Nature Communications* 12.1, p. 7044. DOI: 10.1038/s41467-021-27353-x.

Kirschbaum, D., T. Stanley, and Y. Zhou (2015). "Spatial and Temporal Analysis of a Global Landslide Catalog". *Geomorphology* 249, pp. 4–15. DOI: 10.1016/j.geomorph.2015.03.016.

Kirschbaum, D. B., R. Adler, Y. Hong, S. Hill, and A. Lerner-Lam (2010). "A Global Landslide Catalog for Hazard Applications: Method, Results, and Limitations". *Natural Hazards* 52.3, pp. 561–575. DOI: 10.1007/s11069-009-9401-4.

Kleemans, M. and J. Magruder (2018). "Labour Market Responses to Immigration: Evidence from Internal Migration Driven by Weather Shocks". *The Economic Journal* 128.613, pp. 2032–2065. DOI: 10.1111/ecoj.12510.

Knapp, K. R., M. C. Kruk, D. H. Levinson, H. J. Diamond, and C. J. Neumann (2010). "The International Best Track Archive for Climate Stewardship (IBTrACS): Unifying Tropical Cyclone Data". *Bulletin of the American Meteorological Society* 91.3, pp. 363–376. DOI: 10.1175/2009BAMS2755.1.

Kocornik-Mina, A., T. K. J. McDermott, G. Michaels, and F. Rauch (2020). "Flooded Cities". *American Economic Journal: Applied Economics* 12.2, pp. 35–66. DOI: 10.1257/app.20170066.

Kolstad, C. D. and F. C. Moore (2020). "Estimating the Economic Impacts of Climate Change Using Weather Observations". *Review of Environmental Economics and Policy* 14.1, pp. 1–24. DOI: 10.1093/reep/rez024.

Kugler, M. and H. Rapoport (2007). "International Labor and Capital Flows: Complements or Substitutes?" *Economics Letters* 94.2, pp. 155–162. DOI: 10.1016/j.econlet.2006.06.023.

Kummu, M., M. Kosonen, and S. Masoumzadeh Sayyar (2025). "Downscaled Gridded Global Dataset for Gross Domestic Product (GDP) per Capita PPP over 1990–2022". *Scientific Data* 12.1, p. 178. DOI: 10.1038/s41597-025-04487-x.

Lai, W., Y. Qiu, Q. Tang, C. Xi, and P. Zhang (2023). "The Effects of Temperature on Labor Productivity". *Annual Review of Resource Economics* 15.1, pp. 213–232. DOI: 10.1146/annurev-resource-101222-125630.

Lee, D., W. Anderson, X. Chen, F. Davenport, S. Shukla, R. Sahajpal, M. Budde, J. Rowland, J. Verdin, L. You, M. Ahouangbenon, K. F. Davis, E. Kebede, S. Ehrmann, C. Justice, and C. Meyer (2025). "HarvestStat Africa – Harmonized Subnational Crop Statistics for Sub-Saharan Africa". *Scientific Data* 12.1, p. 690. DOI: 10.1038/s41597-025-05001-z.

Lee, K. and J. Braithwaite (2022). "High-Resolution Poverty Maps in Sub-Saharan Africa". *World Development* 159, p. 106028. DOI: 10.1016/j.worlddev.2022.106028.

Lee, R., C. J. White, M. S. G. Adnan, J. Douglas, M. D. Mahecha, F. E. O'Loughlin, E. Patelli, A. M. Ramos, M. J. Roberts, O. Martius, E. Tubaldi, B. Van Den Hurk, P. J. Ward, and J. Zscheischler (2024). "Reclassifying Historical Disasters: From Single to Multi-Hazards". *Science of The Total Environment* 912, p. 169120. DOI: 10.1016/j.scitotenv.2023.169120.

Lehman, W., M. K. Light, R. Nugent, and J. Burns (2022). *Economic Consequences Assessment Model (ECAM): A Tool & Methodology for Measuring Indirect Economic Effects*. URL: https://www.undrr.org/quick/71614. Pre-published.

Lesk, C., P. Rowhani, and N. Ramankutty (2016). "Influence of Extreme Weather Disasters on Global Crop Production". *Nature* 529.7584, pp. 84–87. DOI: 10.1038/nature16467.

Li, M., S. Cao, Zaichun Zhu, Z. Wang, R. B. Myneni, and S. Piao (2023). *Spatiotemporally Consistent Global Dataset of the GIMMS Normalized Difference Vegetation Index (PKU GIMMS NDVI) from 1982 to 2022 (V1.2)*. Version V1.2. Zenodo. DOI: 10.5281/ZENODO.8253971. (Accessed on 10/21/2024).

Li, M., S. Cao, Z. Zhu, Z. Wang, R. B. Myneni, and S. Piao (2023). "Spatiotemporally Consistent Global Dataset of the GIMMS Normalized Difference Vegetation Index (PKU GIMMS NDVI) from 1982 to 2022". *Earth System Science Data* 15.9, pp. 4181–4203. DOI: 10.5194/essd-15-4181-2023.

Li, X., Y. Zhou, M. Zhao, and X. Zhao (2020). "A Harmonized Global Nighttime Light Dataset 1992–2018". *Scientific Data* 7.1, p. 168. DOI: 10.1038/s41597-020-0510-y.

— (2023). *Harmonization of DMSP and VIIRS Nighttime Light Data from 1992-2021 at the Global Scale*. figshare. DOI: 10.6084/M9.FIGSHARE.9828827.V8. (Accessed on 11/17/2023).

Liu, J., J. Qi, P. Yin, W. Liu, C. He, Y. Gao, L. Zhou, Y. Zhu, H. Kan, R. Chen, and M. Zhou (2024). "Rising Cause-Specific Mortality Risk and Burden of Compound Heatwaves amid Climate Change". *Nature Climate Change* 14.11, pp. 1201–1209. DOI: 10.1038/s41558-024-02137-5.

Liu, L., Y. Wang, and Y. Xu (2024). "A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data". *American Journal of Political Science*, ajps.12723. DOI: 10.1111/ajps.12723.

Liu, L., X. Cao, S. Li, and N. Jie (2024a). "A 31-Year (1990–2020) Global Gridded Population Dataset Generated by Cluster Analysis and Statistical Learning". *Scientific Data* 11.1, p. 124. DOI: 10.1038/s41597-024-02913-0.

— (2024b). *GlobPOP: A 33-Year (1990-2022) Global Gridded Population Dataset Generated by Cluster Analysis and Statistical Learning*. Version 2.0. Zenodo. DOI: 10.5281/zenodo.7813301. (Accessed on 11/03/2024).

Liu, M., Y. Shamdasani, and V. Taraz (2023). "Climate Change and Labor Reallocation: Evidence from Six Decades of the Indian Census". *American Economic Journal: Economic Policy* 15.2, pp. 395–423. DOI: 10.1257/pol.20210129.

Lobell, D. B., W. Schlenker, and J. Costa-Roberts (2011). "Climate Trends and Global Crop Production Since 1980". *Science* 333.6042, pp. 616–620. DOI: 10.1126/science.1204531.

Logothetis, S.-A., V. Salamalikis, A. Gkikas, S. Kazadzis, V. Amiridis, and A. Kazantzidis (2021). "15-Year Variability of Desert Dust Optical Depth on Global and Regional Scales". *Atmospheric Chemistry and Physics* 21.21, pp. 16499–16529. DOI: 10.5194/acp-21-16499-2021.

Magontier, P. and R. Martinez (2023). "Floods \& Urban Density". *SSRN Electronic Journal*. DOI: 10.2139/ssrn.4388712.

Mahajan, P. and D. Yang (2020). "Taken by Storm: Hurricanes, Migrant Networks, and US Immigration". *American Economic Journal: Applied Economics* 12.2, pp. 250–277. DOI: 10.1257/app.20180438.

Mankiw, N. G., D. Romer, and D. N. Weil (1992). "A Contribution to the Empirics of Economic Growth". *The Quarterly Journal of Economics* 107.2, pp. 407–437. DOI: 10.2307/2118477.

Marconcini, M., A. Metz-Marconcini, S. Üreyen, D. Palacios-Lopez, W. Hanke, F. Bachofer, J. Zeidler, T. Esch, N. Gorelick, A. Kakarla, M. Paganini, and E. Strano (2020). "Outlining Where Humans Live, the World Settlement Footprint 2015". *Scientific Data* 7.1, p. 242. DOI: 10.1038/s41597-020-00580-5.

Marí Rivero, I., M. Melchiorri, P. Florio, M. Schiavina, K. Krasnodębska, P. Politis, J. Uhl, M. Pesaresi, L. Maffenini, P. Sulis, M. Crippa, D. Guizzardi, E. Pisoni, C. Belis, D. Oom, A. Branco, D. Mwaniki, E. Kochulem, D. Githira, A. Carioli, D. Ehrlich, P. Tommasi, K. Thomas, and D. Lewis (2024). *GHS Urban Centre Database 2024, Multitemporal and Multidimensional*

Attributes, R2024A. European Commission, Joint Research Centre (JRC). URL: https://data.jrc.ec.europa.eu/dataset/1a338be6-7eaf-480c-9664-3a8ade88cbcd.

Marin-Ferrer, M., K. Poljanšek, and L. Vernaccini (2017). *Index for Risk Management - INFORM :Concept and Methodology, Version 2017*. LU: Publications Office. URL: https://data.europa.eu/doi/10.2760/08037. (Accessed on 07/06/2024).

Meierrieks, D. (2021). "Weather Shocks, Climate Change and Human Health". *World Development* 138, p. 105228. DOI: 10.1016/j.worlddev.2020.105228.

Millock, K. (2015). "Migration and Environment". *Annual Review of Resource Economics* 7.1, pp. 35–60. DOI: 10.1146/annurev-resource-100814-125031.

Minnesota Population Center and Statistics Indonesia (2020). *Integrated Public Use Microdata Series (International: Version 7.3) [Indonesia]*. DOI: 10.18128/D020.V7.3.

Mistry, M. N. (2020). "A High Spatiotemporal Resolution Global Gridded Dataset of Historical Human Discomfort Indices". *Atmosphere* 11.8, p. 835. DOI: 10.3390/atmos11080835.

Moore, M. and D. Wesselbaum (2023). "Climatic Factors as Drivers of Migration: A Review". *Environment, Development and Sustainability* 25.4, pp. 2955–2975. DOI: 10.1007/s10668-022-02191-z.

Mueller, V., G. Sheriff, X. Dou, and C. Gray (2020). "Temporary Migration and Climate Variation in Eastern Africa". *World Development* 126, p. 104704. DOI: 10.1016/j.worlddev.2019.104704.

Muhammad Akhlaq, T. R. Sheltami, and H. T. Mouftah (2012). "A Review of Techniques and Technologies for Sand and Dust Storm Detection". *Reviews in Environmental Science and Bio/Technology* 11.3, pp. 305–322. DOI: 10.1007/s11157-012-9282-y.

Muñoz Sabater, J. (2019). *ERA5-Land Monthly Averaged Data from 1950 to Present*. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). DOI: 10.24381/CDS.68D2BB30. (Accessed on 11/03/2024).

NASA (2023). *Global Landslide Catalog Downloadable Products Gallery*. URL: https://maps.nccs.nasa.gov/arcgis/apps/MapAndAppGallery/index.html?appid=574f26408683485799d02e857e5d9521. (Accessed on 03/27/2023).

National Bureau of Statistics (NBS) (2014). *General Household Survey, Panel 2012-2013, Wave 2*. World Bank, Development Data Group. DOI: 10.48529/KXPY-AA72. (Accessed on 06/10/2025).

National Geophysical Data Center (2024). *Global Historical Tsunami Database*. NOAA National Centers for Environmental Information. DOI: 10.7289/V5PN93H7. (Accessed on 03/27/2023).

Neise, T. and J. R. Diez (2019). "Adapt, Move or Surrender? Manufacturing Firms' Routines and Dynamic Capabilities on Flood Risk Reduction in Coastal Cities of Indonesia". *International Journal of Disaster Risk Reduction* 33, pp. 332–342. DOI: 10.1016/j.ijdrr.2018.10.018.

Neise, T., F. Sohns, M. Breul, and J. R. Diez (2022). "The Effect of Natural Disasters on FDI Attraction: A Sector-Based Analysis over Time and Space". *Natural Hazards* 110.2, pp. 999–1023. DOI: 10.1007/s11069-021-04976-3.

Nickell, S. (1981). "Biases in Dynamic Models with Fixed Effects". *Econometrica* 49.6, p. 1417. DOI: 10.2307/1911408.

Nirandjan, S., E. E. Koks, P. J. Ward, and J. C. Aerts (2021). *A Spatially-Explicit Harmonized Global Dataset of Critical Infrastructure*. Version v1.0.0. Zenodo. DOI: 10.5281/ZENODO.4957647. (Accessed on 06/29/2023).

— (2022). "A Spatially-Explicit Harmonized Global Dataset of Critical Infrastructure". *Scientific Data* 9.1, p. 150. DOI: 10.1038/s41597-022-01218-4.

Niva, V., A. Horton, V. Virkki, M. Heino, M. Kosonen, M. Kallio, P. Kinnunen, G. J. Abel, R. Muttarak, M. Taka, O. Varis, and M. Kummu (2023). "World's Human Migration Patterns in 2000–2019 Unveiled by High-Resolution Data". *Nature Human Behaviour*. DOI: 10.1038/s41562-023-01689-4.

NOAA Climate Prediction Center (2025). *Historical El Nino / La Nina Episodes (1950-Present)*. Cold & Warm Episodes by Season. URL: https://origin.cpc.ncep.noaa.gov/products/analysis_monitoring/ensostuff/ONI_v5.php. (Accessed on 04/07/2025).

NOAA National Centers for Environmental Information (1999). *Global Surface Summary of the Day - GSOD. 1.0*. NOAA National Centers for Environmental Information.

Nordfors, N. (2024). *Droughts and the Growth of Cities*. URL: https://nicklasnordfors.github.io/assets/pdf/droughts_cities.pdf. Pre-published.

Nordhaus, W. (2019). "Climate Change: The Ultimate Challenge for Economics". *American Economic Review* 109.6, pp. 1991–2014. DOI: 10.1257/aer.109.6.1991.

Norling, J. (2022). "Fertility Following Natural Disasters and Epidemics in Africa". *The World Bank Economic Review* 36.4, pp. 955–971. DOI: 10.1093/wber/lhac011.

OCHA's Regional Office for Asia and the Pacific (2021). *Indonesia - Subnational Administrative Boundaries - Humanitarian Data Exchange*. URL: https://data.humdata.org/dataset/indonesia-administrative-boundary-polygons-lines-and-places-levels-0-4b. (Accessed on 04/05/2022).

OECD (2020). *OECD Investment Policy Reviews: Indonesia 2020*. Paris: OECD Publishing. ISBN: 978-92-64-65526-3. DOI: 10.1787/b56512da-en.

OECD and SWAC (2020). *Africapolis*. URL: www.africapolis.org. (Accessed on 12/24/2023).

OECD, United Nations Economic Commission for Africa, and African Development Bank (2022). *Africa's Urbanisation Dynamics 2022: The Economic Power of Africa's Cities*. West African Studies. OECD. ISBN: 978-92-64-91580-0. DOI: 10.1787/3834ed5b-en.

Okou, C., J. Spray, and D. F. Unsal (2022). "Staple Food Prices in Sub-Saharan Africa: An Empirical Assessment". IMF Working Paper Series 2022/135.

Ortiz-Bobea, A. (2021). "The Empirical Analysis of Climate Change Impacts and Adaptation in Agriculture". *Handbook of Agricultural Economics*. Vol. 5. Elsevier, pp. 3981–4073. ISBN: 978-0-323-91501-4. DOI: 10.1016/bs.hesagr.2021.10.002.

Pan, C., Y. Huang, and L. Jin (2024). "Natural Disasters and Corporate Tax Burden: Evidence from Chinese Energy Sector". *Energy Economics* 130, p. 107322. DOI: 10.1016/j.eneco.2024.107322.

Patel, T., S. P. Mullen, and W. R. Santee (2013). "Comparison of Methods for Estimating Wet-Bulb Globe Temperature Index From Standard Meteorological Measurements". *Military Medicine* 178.8, pp. 926–933. DOI: 10.7205/MILMED-D-13-00117.

Pelli, M., J. Tschopp, N. Bezmaternykh, and K. M. Eklou (2023). "In the Eye of the Storm: Firms and Capital Destruction in India". *Journal of Urban Economics* 134, p. 103529. DOI: 10.1016/j.jue.2022.103529.

Pérez-Sindín, X. S., T.-H. K. Chen, and A. V. Prishchepov (2021). "Are Night-Time Lights a Good Proxy of Economic Activity in Rural Areas in Middle and Low-Income Countries? Examining the Empirical Evidence from Colombia". *Remote Sensing Applications: Society and Environment* 24, p. 100647. DOI: 10.1016/j.rsase.2021.100647.

Peri, G. and A. Sasahara (2019). *The Impact of Global Warming on Rural-Urban Migrations: Evidence from Global Big Data*. w25728. Cambridge, MA: National Bureau of Economic Research, w25728. DOI: 10.3386/w25728.

Pesaresi, M. (2023). *GHS-BUILT-S R2023A - GHS Built-up Surface Grid, Derived from Sentinel2 Composite and Landsat, Multitemporal (1975-2030)*. European Commission, Joint Research Centre (JRC). DOI: 10.2905/9F06F36F-4B11-47EC-ABB0-4F8B7B1D72EA. (Accessed on 08/15/2024).

Petrova, K., K. Zantout, S. Zimmermann, V. Niva, M. Kummu, K. Frieler, and J. Schewe (2025). *Differential Effects of Extreme Climate Impacts Onglobal District-Level Net Migration*. DOI: 10.21203/rs.3.rs-6259884/v1. Pre-published.

Pettorelli, N., J. O. Vik, A. Mysterud, J.-M. Gaillard, C. J. Tucker, and N. C. Stenseth (2005). "Using the Satellite-Derived NDVI to Assess Ecological Responses to Environmental Change". *Trends in Ecology & Evolution* 20.9, pp. 503–510. DOI: 10.1016/j.tree.2005.05.011.

Prein, A. F. and G. Holland (2018a). *Daily Large Hail Probability on a Global Scale (1979 to 2015), Version 2, Link to netCDF Files*. text/tab-separated-values. PANGAEA. DOI: 10.1594/PANGAEA.893160. (Accessed on 03/27/2023).

— (2018b). "Global Estimates of Damaging Hail Hazard". *Weather and Climate Extremes* 22, pp. 10–23. DOI: 10.1016/j.wace.2018.10.004.

Pronk, M., A. Hooijer, D. Eilander, A. Haag, T. De Jong, M. Vousdoukas, R. Vernimmen, H. Ledoux, and M. Eleveld (2024). "DeltaDTM: A Global Coastal Digital Terrain Model". *Scientific Data* 11.1, p. 273. DOI: 10.1038/s41597-024-03091-9.

Provenzano, S. (2024). "Accountability Failure in Isolated Areas: The Cost of Remoteness from the Capital City". *Journal of Development Economics* 167, p. 103214. DOI: 10.1016/j.jdeveco.2023.103214.

Pullabhotla, H. K., M. Zahid, S. Heft-Neal, V. Rathi, and M. Burke (2023). "Global Biomass Fires and Infant Mortality". *Proceedings of the National Academy of Sciences* 120.23, e2218210120. DOI: 10.1073/pnas.2218210120.

PwC (2024). *Indonesia: Corporate - Taxes on Corporate Income*. Worldwide Tax Summaries. URL: https://taxsummaries.pwc.com/indonesia/corporate/taxes-on-corporate-income. (Accessed on 05/15/2025).

Raleigh, C., R. Kishi, and A. Linke (2023). "Political Instability Patterns Are Obscured by Conflict Dataset Scope Conditions, Sources, and Coding Choices". *Humanities and Social Sciences Communications* 10.1, p. 74. DOI: 10.1057/s41599-023-01559-4.

Rambachan, A. and J. Roth (2023). "A More Credible Approach to Parallel Trends". *Review of Economic Studies* 90.5, pp. 2555–2591. DOI: 10.1093/restud/rdad018.

Reed, B. C., J. F. Brown, D. VanderZee, T. R. Loveland, J. W. Merchant, and D. O. Ohlen (1994). "Measuring Phenological Variability from Satellite Imagery". *Journal of Vegetation Science* 5.5, pp. 703–714. DOI: 10.2307/3235884.

Reinhardt, R. (2024a). *A Global Gridded Data Collection on Climate Impact Drivers and Economics from 2000-2019*. Version 1.0. DOI: 10.5281/zenodo.14220525.

— (2024b). *Replication Package for Platform for Economic (Climate) Hazard (PEACH)*. Version 1.0. Zenodo. DOI: 10.5281/ZENODO.14220364.

— (2024c). "Shaking up Foreign Finance: FDI in a Post-Disaster World". *Economics of Disasters and Climate Change* 8.2, pp. 317–348. DOI: 10.1007/s41885-024-00148-2.

Ren, H., X. Huang, J. Yang, and G. Zhou (2025). "Improving 30-Meter Global Impervious Surface Area (GISA) Mapping: New Method and Dataset". *ISPRS Journal of Photogrammetry and Remote Sensing* 220, pp. 354–376. DOI: 10.1016/j.isprsjprs.2024.12.023.

Rentschler, J. (2024). *Global Flood Exposure: Gridded Exposure Headcounts by Country | Data Catalog*. URL: https://datacatalog.worldbank.org/search/dataset/0062763/Global-Flood-Exposure--Gridded-exposure-headcounts-by-country. (Accessed on 10/18/2024).

Rentschler, J., P. Avner, M. Marconcini, R. Su, E. Strano, M. Vousdoukas, and S. Hallegatte (2023). "Global Evidence of Rapid Urban Growth in Flood Zones since 1985". *Nature* 622.7981, pp. 87–92. DOI: 10.1038/s41586-023-06468-9.

Rentschler, J., E. Kim, S. Thies, S. De, V. Robbe, A. Erman, and S. Hallegatte (2021). *Floods and Their Impacts on Firms Evidence from Tanzania*. 9774. The World Bank. URL: https://openknowledge.worldbank.org/handle/10986/36282.

Rentschler, J., M. Salhab, and B. A. Jafino (2022). "Flood Exposure and Poverty in 188 Countries". *Nature Communications* 13.1, p. 3527. DOI: 10.1038/s41467-022-30727-4.

Ridder, N. N., A. J. Pitman, S. Westra, A. Ukkola, H. X. Do, M. Bador, A. L. Hirsch, J. P. Evans, A. Di Luca, and J. Zscheischler (2020). "Global Hotspots for the Occurrence of Compound Events". *Nature Communications* 11.1, p. 5956. DOI: 10.1038/s41467-020-19639-3.

Rockström, J., W. Steffen, K. Noone, Å. Persson, F. S. Chapin, E. F. Lambin, T. M. Lenton, M. Scheffer, C. Folke, H. J. Schellnhuber, B. Nykvist, C. A. De Wit, T. Hughes, S. Van Der Leeuw, H. Rodhe, S. Sörlin, P. K. Snyder, R. Costanza, U. Svedin, M. Falkenmark, L. Karlberg, R. W. Corell, V. J. Fabry, J. Hansen, B. Walker, D. Liverman, K. Richardson, P. Crutzen, and J. A. Foley (2009). "A Safe Operating Space for Humanity". *Nature* 461.7263, pp. 472–475. DOI: 10.1038/461472a.

Rossi-Hansberg, E. and J. Zhang (2025). *Local GDP Estimates Around the World*. w33458. Cambridge, MA: National Bureau of Economic Research, w33458. DOI: 10.3386/w33458.

Rosvold, E. and H. Buhaug (2021a). *Geocoded Disasters (GDIS) Dataset*. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC). DOI: 10.7927/ZZ3B-8Y61. (Accessed on 02/27/2023).

Rosvold, E. and H. Buhaug (2021b). "GDIS, a Global Dataset of Geocoded Disaster Locations". *Scientific Data* 8.1, p. 61. DOI: 10.1038/s41597-021-00846-6.

Ruane, A. C., R. Vautard, R. Ranasinghe, J. Sillmann, E. Coppola, N. Arnell, F. A. Cruz, S. Dessai, C. E. Iles, A. K. M. S. Islam, R. G. Jones, M. Rahimi, D. R. Carrascal, S. I. Seneviratne, J. Servonnat, A. A. Sörensson, M. B. Sylla, C. Tebaldi, W. Wang, and R. Zaaboul (2022). "The Climatic Impact-Driver Framework for Assessment of Risk-Relevant Climate Information". *Earth's Future* 10.11, e2022EF002803. DOI: 10.1029/2022EF002803.

Runfola, D., A. Anderson, H. Baier, M. Crittenden, E. Dowker, S. Fuhrig, S. Goodman, G. Grimsley, R. Layko, G. Melville, M. Mulder, R. Oberman, J. Panganiban, A. Peck, L. Seitz, S. Shea, H. Slevin, R. Youngerman, and L. Hobbs (2020). "geoBoundaries: A Global Database of Political Administrative Boundaries". *PLOS ONE* 15.4. Ed. by W. Tang, e0231866. DOI: 10.1371/journal.pone.0231866.

Sadiddin, A., A. Cattaneo, M. Cirillo, and M. Miller (2019). "Food Insecurity as a Determinant of International Migration: Evidence from Sub-Saharan Africa". *Food Security* 11.3, pp. 515–530. DOI: 10.1007/s12571-019-00927-w.

Šakić Trogrlić, R., H. E. Thompson, E. Y. Menteşe, E. Hussain, J. C. Gill, F. E. Taylor, E. Mwangi, E. Öner, V. G. Bukachi, and B. D. Malamud (2024). "Multi-Hazard Interrelationships and Risk Scenarios in Urban Areas: A Case of Nairobi and Istanbul". *Earth's Future* 12.9, e2023EF004413. DOI: 10.1029/2023EF004413.

Sant'Anna, P. H. C. and J. Zhao (2020). "Doubly Robust Difference-in-Differences Estimators". *Journal of Econometrics* 219.1, pp. 101–122. DOI: 10.1016/j.jeconom.2020.06.003.

Santini, M., S. Noce, M. Antonelli, and L. Caporaso (2022). "Complex Drought Patterns Robustly Explain Global Yield Loss for Major Crops". *Scientific Reports* 12.1, p. 5792. DOI: 10.1038/s41598-022-09611-0.

Sauter, C., J. L. Catto, H. J. Fowler, S. Westra, and C. J. White (2023). "Compounding Heatwave-Extreme Rainfall Events Driven by Fronts, High Moisture, and Atmospheric Instability". *Journal of Geophysical Research: Atmospheres* 128.21, e2023JD038761. DOI: 10.1029/2023JD038761.

Schwarzwald, K. and K. Geil (2024). "Xagg: A Python Package to Aggregate Gridded Data onto Polygons". *Journal of Open Source Software* 9.104, p. 7239. DOI: 10.21105/joss.07239.

SEDAC-CIESIN (2017). *Gridded Population of the World, Version 4 (GPWv4): Population Density, Revision 11*. Palisades, NY: Socioeconomic Data and Applications Center (SEDAC). DOI: 10.7927/H49C6VHW.

Sen, S., N. C. Nayak, and W. K. Mohanty (2022). "Estimating Household Vulnerability to Tropical Cyclones: An Investigation of Tropical Cyclone Shocks in Coastal Villages of Eastern India". *International Journal of Disaster Risk Reduction* 83, p. 103404. DOI: 10.1016/j.ijdrr.2022.103404.

Seneviratne, S., N. Nicholls, D. Easterling, C. Goodess, S. Kanae, J. Kossin, Y. Luo, J. Marengo, K. McInnes, M. Rahimi, M. Reichstein, A. Sorteberg, C. Vera, and X. Zhang (2012). "Changes in Climate Extremes and Their Impacts on the Natural Physical Environmen". *Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation*. Ed. by C. Field, V. Barros, T. Stocker, D. Qin, D. Dokken, K. Ebi, M. Mastrandrea, K. Mach, G.-K. Plattner, S. Allen, M. Tignor, and P. Midgley. Vol. Special Report of Working Groups I and II of the Intergovernmental Panel on Climate Change. Cambridge, UK and New York, USA: Cambridge University Press, pp. 109–230.

Seneviratne, S., X. Zhang, M. Adnan, W. Badi, C. Dereczynski, A. Di Luca, S. Ghosh, I. Iskandar, J. Kossin, S. Lewis, F. Otto, I. Pinto, M. Satoh, S. Vicente-Serrano, M. Wehner, and B. Zhou (2021). "Chapter 11 - Weather and Climate Extreme Events in a Changing Climate". *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by V. Masson-Delmotte, P. Zhai, A. Pirani, S. Connors, C. Pean, S. Berger, N. Chaud, Y. Chen, L. Goldfarb, M. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. Matthews, T. Maycook, T. Waterfield, O. Yelekci, R. Yu, and B. Zhou. Cambridge, United Kingdom and New York, NY: Cambridge University Press, pp. 1513–1766. DOI: 10.1017/9781009157896.013.

Sepúlveda, I., J. S. Haase, P. L.-F. Liu, M. Grigoriu, and P. Winckler (2021). "Non-Stationary Probabilistic Tsunami Hazard Assessments Incorporating Climate-Change-Driven Sea Level Rise". *Earth's Future* 9.6. DOI: 10.1029/2021EF002007.

Shabnam, N. (2014). "Natural Disasters and Economic Growth: A Review". *International Journal of Disaster Risk Science* 5.2, pp. 157–163. DOI: 10.1007/s13753-014-0022-5.

Shan, L. and S. X. Gong (2012). "Investor Sentiment and Stock Returns: Wenchuan Earthquake". *Finance Research Letters* 9.1, pp. 36–47. DOI: 10.1016/j.frl.2011.07.002.

Shen, S., C. Li, A. Van Donkelaar, N. Jacobs, C. Wang, and R. V. Martin (2024). "Enhancing Global Estimation of Fine Particulate Matter Concentrations by Including Geophysical a Priori Information in Deep Learning". *ACS ES&T Air* 1.5, pp. 332–345. DOI: 10.1021/acsestair.3c00054.

Sherman, L., J. Proctor, H. Druckenmiller, H. Tapia, and S. Hsiang (2023). *Global High-Resolution Estimates of the United Nations Human Development Index Using Satellite Imagery and Machine-*

learning. w31044. Cambridge, MA: National Bureau of Economic Research, w31044. DOI: 10.3386/w31044.

Steenbergen, V., S. Hebous, M. M. Wihardja, and A. T. Pradana (2020). *The Effect of FDI on Indonesia's Jobs, Wages, and Structural Transformation*. World Bank. DOI: 10.1596/36188.

Steer, P., L. Jeandet, N. Cubas, O. Marc, P. Meunier, M. Simoes, R. Cattin, J. B. H. Shyu, M. Mouyen, W.-T. Liang, T. Theunissen, S.-H. Chiang, and N. Hovius (2020). "Earthquake Statistics Changed by Typhoon-Driven Erosion". *Scientific Reports* 10.1, p. 10899. DOI: 10.1038/s41598-020-67865-y.

Stokey, N. L. (2016). "Wait-and-See: Investment Options under Policy Uncertainty". *Review of Economic Dynamics* 21, pp. 246–265. DOI: 10.1016/j.red.2015.06.001.

Sun, L. and S. Abraham (2021). "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects". *Journal of Econometrics* 225.2, pp. 175–199. DOI: 10.1016/j.jeconom.2020.09.006.

Sundberg, R. and E. Melander (2013). "Introducing the UCDP Georeferenced Event Dataset". *Journal of Peace Research* 50.4, pp. 523–532. DOI: 10.1177/0022343313484347.

Sutton, P. C., C. D. Elvidge, and T. Ghosh (2007). "Estimation of Gross Domestic Product at Sub-national Scales Using Nighttime Satellite Imagery." *International Journal of Ecological Economics & Statistics* 8.S07, pp. 5–21.

Suyanto, R. A. Salim, and H. Bloch (2009). "Does Foreign Direct Investment Lead to Productivity Spillovers? Firm Level Evidence from Indonesia". *World Development* 37.12, pp. 1861–1876. DOI: 10.1016/j.worlddev.2009.05.009.

Tao, S., Z. Ao, J.-P. Wigneron, S. Saatchi, P. Ciais, J. Chave, T. Le Toan, P.-L. Frison, X. Hu, C. Chen, L. Fan, M. Wang, J. Zhu, X. Zhao, X. Li, X. Liu, Y. Su, T. Hu, Q. Guo, Z. Wang, Z. Tang, Y. Y. Liu, and J. Fang (2023). "A Global Long-Term, High-Resolution Satellite Radar Backscatter Data Record (1992–2022+): Merging C-band ERS/ASCAT and Ku-band QSCAT". *Earth System Science Data* 15.4, pp. 1577–1596. DOI: 10.5194/essd-15-1577-2023.

Taszarek, M., J. T. Allen, M. Marchio, and H. E. Brooks (2021). "Global Climatology and Trends in Convective Environments from ERA5 and Rawinsonde Data". *npj Climate and Atmospheric Science* 4.1, p. 35. DOI: 10.1038/s41612-021-00190-x.

Thalheimer, L., N. Choquette-Levy, and F. Garip (2022). "Compound Impacts from Droughts and Structural Vulnerability on Human Mobility". *iScience* 25.12, p. 105491. DOI: 10.1016/j.isci.2022.105491.

Thalheimer, L., F. Otto, and S. Abele (2021). "Deciphering Impacts and Human Responses to a Changing Climate in East Africa". *Frontiers in Climate* 3, p. 692114. DOI: 10.3389/fclim.2021.692114.

The World Bank (2021a). *How Indonesia Strengthened Its Disaster Response with Risk Finance and Insurance*. URL: https://www.worldbank.org/en/news/feature/2021/11/17/how-indonesia-strengthened-its-disaster-response-with-risk-finance-and-insurance. (Accessed on 05/15/2025).

— (2021b). *Indonesia - Disaster Risk Finance and Insurance Project (English)*. Washington, D.C.: World Bank Group. URL: http://documents.worldbank.org/curated/en/316601611543685552.

— (2022a). *Foreign Direct Investment, Net Inflows (% of GDP) - Indonesia*. URL: https://data.worldbank.org/indicator/BX.KLT.DINV.WD.GD.ZS?end=2020&locations=ID&start=1970&view=chart. (Accessed on 02/27/2022).

— (2022b). *Indonesia Database for Policy and Economic Research (INDO-DAPOER)*. URL: https://datacatalog.worldbank.org/dataset/indonesia-database-policy-and-economic-research. (Accessed on 02/27/2022).

— (2024). *GDP (Current USD)*. URL: https://data.worldbank.org/indicator/NY.GDP.MKTP.CD. (Accessed on 09/21/2024).

Tol, R. S. J. (2018). "The Economic Impacts of Climate Change". *Review of Environmental Economics and Policy* 12.1, pp. 4–25. DOI: 10.1093/reep/rex027.

Tuholske, C., K. Caylor, C. Funk, A. Verdin, S. Sweeney, K. Grace, P. Peterson, and T. Evans (2021). "Global Urban Population Exposure to Extreme Heat". *Proceedings of the National Academy of Sciences* 118.41, e2024792118. DOI: 10.1073/pnas.2024792118.

Tziokas, N., C. Zhang, G. C. Drolias, and P. M. Atkinson (2023). "Downscaling Satellite Night-Time Lights Imagery to Support within-City Applications Using a Spatially Non-Stationary Model". *International Journal of Applied Earth Observation and Geoinformation* 122, p. 103395. DOI: 10.1016/j.jag.2023.103395.

U.S. Environmental Protection Agency (2024). *Final Reconsideration of the National Ambient Air Quality Standards for Particulate Matter (PM)*. URL: https://www.epa.gov/system/files/documents/2024-02/pm-naaqs-air-quality-index-fact-sheet.pdf. (Accessed on 11/13/2024).

UN Economic and Social Affairs (Population Division) (2019). *International Migrant Stock 2019*. URL: https://www.un.org/development/desa/pd/content/international-migrant-stock. (Accessed on 12/04/2021).

UN OCHA Regional Office for Asia Pacific (OCHA-ROAP) (2011). *Natural Hazard Risks Indonesia*. URL: https://reliefweb.int/attachments/c919bdef-6a8d-35e6-9149-502a43170679/map.pdf. (Accessed on 10/10/2020).

UNCTAD (2017). *World Investment Report 2017: Investment and the Digital Economy*. United Nations Conference on Trade and Development. 52 pp. ISBN: 978-92-1-112911-3.

— (2021). *Foreign Direct Investment: Inward and Outward Flows and Stock, Annual*. URL: https://unctadstat.unctad.org/. (Accessed on 08/09/2021).

UNDESA (2019). *World Urbanization Prospects: The 2018 Revision Produced by the Population Division of the UN Department of Economic and Social Affairs*. New York: United Nations.

UNDP Indonesia and BNPB Indonesia (2019). *Data Dan Informasi Bencana Indonesia - Lessons Learned DiBi*. URL: https://bnpb.go.id/uploads/migration/pubs/442.pdf. (Accessed on 03/22/2021).

UNDRR (2017a). *Flood Hazard 200 Years*. URL: https://risk.preventionweb.net/index.html. (Accessed on 01/19/2021).

— (2017b). *Peak Ground Acceleration PGA 250 Years*. URL: https://risk.preventionweb.net/capraviewer/. (Accessed on 01/19/2021).

— (2022a). *Definition of Basic Effects*. URL: https://www.desinventar.net/effects.html. (Accessed on 10/08/2022).

— (2022b). *Desinventar Sendai - Framework for Disaster Risk Reduction*. URL: https://www.desinventar.net/DesInventar/. (Accessed on 02/12/2022).

UNEP/DEWA/GRID-Europe and P. Peduzzi (2014). *Fires - Density*. URL: https://web.archive.org/web/20220306050958/https://preview.grid.unep.ch/index.php?preview=data&events=fires&evcat=3&lang=eng. (Accessed on 01/19/2021).

UNHCR (2024). *Informal Settlements*. Emergency Handbook. URL: https://emergency.unhcr.org/emergency-assistance/shelter-camp-and-settlement/settlements/informal-settlements. (Accessed on 05/11/2025).

United Nations Office for Disaster Risk Reduction (UNDRR) (2024). *Desinventar Sendai - Framework for Disaster Risk Reduction*. URL: https://www.desinventar.net/. (Accessed on 08/30/2024).

United States Geological Survey (USGS) (2017). *ShakeMap - Earthquake Ground Motion and Shaking Intensity Maps*. U.S. Geological Survey. DOI: 10.5066/F7W957B2. (Accessed on 10/18/2024).

— (2022). *The Modified Mercalli Intensity Scale*. URL: https://www.usgs.gov/programs/earthquake-hazards/modified-mercalli-intensity-scale. (Accessed on 02/12/2022).

— (2023). *ANSS Comprehensive Earthquake Catalog (ComCat)*. URL: https://earthquake.usgs.gov/data/comcat/index.php#6. (Accessed on 06/29/2023).

UNSD (2023). *Standard Country or Area Codes for Statistical Use (M49)*. URL: https://unstats.un.org/unsd/methodology/m49/. (Accessed on 05/03/2023).

US Geological Survey (2020). *ANSS Comprehensive Earthquake Catalog (ComCat)*.

Van der Borght, R., O. A. Ishizawa, J. Thuret, and J. Muñoz (2025). *Dry Spells, Urban Swells: Analyzing the Drought-Induced Expansion of Cities in Sub-Saharan Africa*. 11099. World Bank, Washington, DC: World Bank. URL: http://hdl.handle.net/10986/43051.

Vermote, E. and R. Wolfe (2015). *MOD09GQ MODIS/Terra Surface Reflectance Daily L2G Global 250m SIN Grid V006*. NASA EOSDIS Land Processes Distributed Active Archive Center. DOI: 10.5067/MODIS/MOD09GQ.006. (Accessed on 10/18/2024).

Vigdor, J. (2008). "The Economic Aftermath of Hurricane Katrina". *Journal of Economic Perspectives* 22.4, pp. 135–154. DOI: 10.1257/jep.22.4.135.

Vitolo, C., F. Di Giuseppe, B. Krzeminski, and J. San-Miguel-Ayanz (2019). "A 1980–2018 Global Fire Danger Re-Analysis Dataset for the Canadian Fire Weather Indices". *Scientific Data* 6.1, p. 190032. DOI: 10.1038/sdata.2019.32.

Vousdoukas, M. (2023). *Global Coastal Flood Maps*. Zenodo. DOI: 10.5281/ZENODO.8057902. (Accessed on 04/02/2025).

Vousdoukas, M. I., L. Mentaschi, E. Voukouvalas, M. Verlaan, S. Jevrejeva, L. P. Jackson, and L. Feyen (2018). "Global Probabilistic Projections of Extreme Sea Levels Show Intensification of Coastal Flood Hazard". *Nature Communications* 9.1, p. 2360. DOI: 10.1038/s41467-018-04692-w.

Wald, D. J., V. Quitoriano, T. H. Heaton, H. Kanamori, C. W. Scrivner, and C. B. Worden (1999). *TriNet "ShakeMaps": Rapid Generation of Peak Ground Motion and Intensity Maps for Earthquakes in Southern California*. 3, pp. 537–554.

Watts, N., W. N. Adger, S. Ayeb-Karlsson, Y. Bai, P. Byass, D. Campbell-Lendrum, T. Colbourn, P. Cox, M. Davies, M. Depledge, A. Depoux, P. Dominguez-Salas, P. Drummond, P. Ekins, A. Flahault, D. Grace, H. Graham, A. Haines, I. Hamilton, A. Johnson, I. Kelman, S. Kovats, L. Liang, M. Lott, R. Lowe, Y. Luo, G. Mace, M. Maslin, K. Morrissey, K. Murray, T. Neville, M. Nilsson, T. Oreszczyn, C. Parthemore, D. Pencheon, E. Robinson, S. Schütte, J. Shumake-Guillemot, P. Vineis, P. Wilkinson, N. Wheeler, B. Xu, J. Yang, Y. Yin, C. Yu, P. Gong, H. Montgomery, and A. Costello (2017). "The Lancet Countdown: Tracking Progress on Health

and Climate Change". *The Lancet* 389.10074, pp. 1151–1164. DOI: 10.1016/S0140-6736(16)32124-9.

Winkler, K., R. Fuchs, M. Rounsevell, and M. Herold (2020). *HILDA+ Global Land Use Change between 1960 and 2019*. text/tab-separated-values. PANGAEA. DOI: 10.1594/PANGAEA.921846. (Accessed on 06/29/2023).

— (2021). "Global Land Use Changes Are Four Times Greater than Previously Estimated". *Nature Communications* 12.1, p. 2501. DOI: 10.1038/s41467-021-22702-2.

WMO (2024). *WMO Climatological Normals | World Meteorological Organization*. URL: https://community.wmo.int/en/wmo-climatological-normals. (Accessed on 03/25/2024).

Wooldridge, J. M. (2021). "Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators". *SSRN Electronic Journal*. DOI: 10.2139/ssrn.3906345.

WorldPop (2018). *Global 1km Age/Sex Structures*. University of Southampton. DOI: 10.5258/SOTON/WP00654. (Accessed on 11/15/2024).

— (2020). *Global 100m Population Total Adjusted to Match the Corresponding UNPD Estimate*. University of Southampton. DOI: 10.5258/SOTON/WP00660. (Accessed on 11/21/2024).

Wu, J., H. Yao, X. Chen, G. Wang, X. Bai, and D. Zhang (2022). "A Framework for Assessing Compound Drought Events from a Drought Propagation Perspective". *Journal of Hydrology* 604, p. 127228. DOI: 10.1016/j.jhydrol.2021.127228.

Xu, Y. (2023). "Causal Inference with Time-Series Cross-Sectional Data: A Reflection". *Oxford Handbook of Engaged Methodological Pluralism in Political Science*. Ed. by J. M. Box-Steffensmeier, D. P. Christenson, and V. Sinclair-Chapman. 1st ed. Oxford University Press. ISBN: 978-0-19-286828-2 978-0-19-196422-0. DOI: 10.1093/oxfordhb/9780192868282.013.30.

Xu, Y. and L. Licheng (2018). *panelView: An R Package of Visualizing Panel Data*. URL: https://yiqingxu.org/packages/panelview/index.html.

Zhang, H., G. Dong, B. Li, Z. Xie, C. Miao, F. Yang, Y. Gao, X. Meng, D. Yang, Y. Liu, H. Zhang, L. Wu, F. Shi, Y. Chen, W. Wu, E. Laszkiewicz, Y. Liang, B. Lu, J. Yao, and X. Li (2024). "Developing an Annual Global Sub-National Scale Economic Data from 1992 to 2021 Using Nighttime Lights and Deep Learning". *International Journal of Applied Earth Observation and Geoinformation* 133, p. 104086. DOI: 10.1016/j.jag.2024.104086.

Zhao, M., J. K. W. Lee, T. Kjellstrom, and W. Cai (2021). "Assessment of the Economic Impact of Heat-Related Labor Productivity Loss: A Systematic Review". *Climatic Change* 167.1-2, p. 22. DOI: 10.1007/s10584-021-03160-7.

Zhu, H., L. Duan, Y. Guo, and K. Yu (2016). "The Effects of FDI, Economic Growth and Energy Consumption on Carbon Emissions in ASEAN-5: Evidence from Panel Quantile Regression". *Economic Modelling* 58, pp. 237–248. DOI: 10.1016/j.econmod.2016.05.003.

Zscheischler, J., O. Martius, S. Westra, E. Bevacqua, C. Raymond, R. M. Horton, B. Van Den Hurk, A. AghaKouchak, A. Jézéquel, M. D. Mahecha, D. Maraun, A. M. Ramos, N. N. Ridder, W. Thiery, and E. Vignotto (2020). "A Typology of Compound Weather and Climate Events". *Nature Reviews Earth & Environment* 1.7, pp. 333–347. DOI: 10.1038/s43017-020-0060-z.

Zscheischler, J., S. Westra, B. J. J. M. Van Den Hurk, S. I. Seneviratne, P. J. Ward, A. Pitman, A. AghaKouchak, D. N. Bresch, M. Leonard, T. Wahl, and X. Zhang (2018). "Future Climate Risk from Compound Events". *Nature Climate Change* 8.6, pp. 469–477. DOI: 10.1038/s41558-018-0156-3.

Zuzak, C., E. Goodenough, C. Stanton, M. Mowrer, A. Sheenan, B. Roberts, P. McGuire, and J. Rozelle (2023). *National Risk Index Technical Documentation*. Washington, DC, p. 272. URL: https://www.fema.gov/sites/default/files/documents/fema_national-risk-index_technical-documentation.pdf.