



**Université Paris 1 Panthéon-Sorbonne**  
**École doctorale d'Économie (ED 465)**  
**Centre d'Économie de la Sorbonne (CES)**

## **Ph.D. Thesis**

Submitted for the Degree of Doctor of Philosophy in Economics

*Prepared and publicly defended on January 21, 2025 by*  
**Jean-Michel Dagba**

### *Title*

# **Playing favourites: Investigating Information Search Before Set Selection**

*Under the supervision of*  
**Jean-Christophe Vergnaud & Vincent de Gardelle**

*With the support of*  
**Sandrine Bélier (Humans Matter company)**

### **Jury**

- |   |                  |
|---|------------------|
| = Angela Sutan, <i>professeure des universités</i> , ESSEC Business School          | <i>Reviewer</i>  |
| = Alexandre Zénon, <i>chercheur CNRS</i> , Université de Bordeaux                   | <i>Reviewer</i>  |
| = Antoine Mandel, <i>professeur des universités</i> , Université Paris 1            | <i>President</i> |
| = Wim de Neys, <i>directeur de recherche CNRS</i> , Université Paris Cité           | <i>Examiner</i>  |
| = Vincent de Gardelle, <i>directeur de recherche CNRS</i> , Université Paris 1      | <i>Advisor</i>   |
| = Jean-Christophe Vergnaud, <i>directeur de recherche CNRS</i> , Université Paris 1 | <i>Advisor</i>   |

## Résumé

Dans de nombreuses situations, il est nécessaire de choisir un ensemble d'options plutôt qu'une seule. Par exemple, une étudiante qui sélectionne ses cours ou un recruteur choisissant les candidats à recevoir en entretien. Ces choix impliquent une recherche d'informations : l'étudiante se renseigne sur les cours, et le recruteur examine CV et lettres de motivation.

Ce processus est complexe car il nécessite de comparer des combinaisons d'options, et non des options isolées. Les informations les plus utiles à explorer ne concernent pas toujours les options favorites, celles-ci étant généralement incluses dans la sélection finale. Ce type de recherche est peu étudié, malgré sa complexité et ses applications. Cette thèse examine les stratégies d'exploration d'informations pour sélectionner un ensemble d'options. Nous cherchons à comprendre comment les individus recherchent des informations, quels facteurs influencent leurs stratégies, et si des stratégies optimales peuvent leur être enseignées.

Les recherches montrent que les individus, face à une sélection unique, sont sous-optimaux par rapport aux modèles bayésiens. Ils utilisent des heuristiques, des stratégies simplifiées influencées par leurs biais cognitifs, pour réduire la complexité. Nous supposons que ces limites de rationalité influencent également la recherche d'informations avant une sélection d'ensemble.

Quatre séries d'expériences ont été menées. Dans un paradigme original, les participants devaient choisir le meilleur ensemble d'options après avoir réalisé une recherche d'informations limitée. L'objectif est de trouver les 5 meilleures candidatures parmi un ensemble de 10. Pour cela, un participant a accès un score A pour chaque candidature. Il doit décider pour quelles candidatures il souhaite consulter le score B avant de choisir les 5 candidatures parmi les 10 qui ont le meilleur score total (score A + score B). Les résultats révèlent deux biais majeurs : (1) Biais d'exploration : privilégier les options favorites lors de la recherche d'informations, alors que l'effort devrait se concentrer sur les options incertaines. (2) Biais de sélection : favoriser les options explorées, même si des options non-explorées seraient des meilleurs choix.

Le chapitre 2 montre que les biais ne sont pas fortement corrélés à d'autres biais cognitifs ou à l'expertise en recrutement. Le chapitre 3 indique que les individus surestiment la qualité de leurs stratégies. Enfin, le chapitre 4 révèle que les participants peuvent apprendre des stratégies optimales.

# Abstract

In many situations, it is necessary to select a set of options rather than a single one. For example, a student choosing their courses or a recruiter selecting candidates to invite for interviews. These decisions involve an information-gathering process: the student investigates the course content, while the recruiter reviews CVs and cover letters.

This process is complex as it requires comparing combinations of options rather than evaluating options in isolation. The most useful information to explore does not always concern the favourite options, as these are generally included in the final selection. This type of information search is overlooked despite its complexity and practical applications. This thesis investigates information search strategies for selecting a set of options. We aim to understand how individuals search for information, what factors influence their strategies, and whether optimal strategies can be taught.

Research shows that individuals are suboptimal compared to Bayesian models when conducting information searches for single-option selection. They rely on heuristics—simplified strategies influenced by cognitive biases—to reduce complexity. We hypothesise that these limitations of rationality also affect information searches preceding set selection.

Four experimental studies were conducted. In an original paradigm, participants were tasked with selecting the best set of options after conducting a limited information search. The goal was to identify the five best candidates out of a set of ten. Participants initially had access to a score A for each candidate and decided for which candidates they wanted to examine the score B before selecting the five candidates with the highest combined scores (A + B). The results revealed two major biases: (1) Exploration bias: prioritising favourite options during the information search, whereas efforts should focus on uncertain options. (2) Selection bias: favouring explored options even when unexplored options would yield better choices.

Chapter 2 demonstrates that these biases are not strongly correlated with other cognitive biases or recruitment expertise. Chapter 3 shows that individuals overestimate the quality of their strategies. Finally, Chapter 4 reveals that participants can learn and apply optimal strategies.

## Mots clés

Recherche d'information – Prise de décision – Sélection d'ensemble – Heuristiques  
Biais cognitifs – Surconfiance – Biais d'exploration – Biais de sélection  
Cohérence interne – Remédiation – Incitation financières – Économie comportementale  
Économie expérimentale – Psychologie cognitive – Psychologie expérimentale

## Keywords

Information search – Decision making – Set selection – Heuristics  
Cognitive bias – Overconfidence – Exploration bias – Selection bias Internal consistency  
Remediation – Financial incentives – Behavioural economics – Experimental economics  
Cognitive psychology – Experimental psychology

## Remerciements

Tout d'abord, merci Jean-Christophe & Vincent de m'avoir donné ma chance et de m'avoir fait confiance. On dit souvent, entre doctorants, que la direction de thèse joue un rôle crucial, non seulement dans la conduite des travaux, mais surtout sur la santé mentale des doctorants. Je pense que j'ai pu mener à bien cette thèse en étant bien dans ma thèse grâce à vous. Je souhaite à toute personne qui souhaite faire une thèse d'être encadré comme je l'ai été. Je suis content de la confiance et de la relation que nous avons réussi à développer, bien au-delà de la thèse. Merci beaucoup !

Merci à Angela Sutan, Alexandre Zénon, Wim de Neys et Antoine Mandel d'avoir accepté de faire partie de mon jury. J'espère que la lecture de ma thèse pendant les fêtes de fin d'année pourra vous servir d'excuses pour éviter des situations embarrassantes... C'est le seul avantage que j'ai trouvé ; désolé pour ce timing peu commode ! Merci aussi à Vincent Berthet, Jérôme Sackur, Isabelle Régner, Sandrine Bélier et Thomas Breda d'avoir fait partie de mes comités de thèse. Vos conseils et votre sympathie à l'égard de mes travaux m'ont beaucoup aidé à aboutir cette thèse. En particulier merci à Vincent qui m'a inspiré et m'a donné envie de m'intéresser aux biais cognitifs à la rationalité limitée quand je n'étais encore qu'en licence à Nancy.

Merci à toutes les personnes qui m'ont aidé à construire mon chemin professionnel jusqu'à l'obtention de cette thèse. En particulier Tania et Mohamed de L'École du Recrutement qui m'ont aidé à trouver des recruteurs pour des expériences ; Julie Grèzes et Rocco Mennella qui m'ont permis de faire mon premier stage en labo ; Cog'X qui m'a permis de découvrir les sciences cognitives en entreprise ; Aurélien & Hanaë qui m'ont permis de présenter ma thèse lors d'un TEDx.

Merci à vous, les doctorants, auprès de qui j'ai trouvé ma place. Rien contre les gens qui ont choisi de faire autre chose de leur vie hein, mais il y a un truc différent entre nous. J'ai d'abord pensé que venir au labo n'était utile que pour les réunions avec mes directeurs. Quelle erreur ! Peu à peu, j'ai compris la chance que j'avais d'être avec vous, d'apprendre à vous connaître, de pouvoir compter sur votre soutien et que l'on puisse s'entraider. Merci en premier lieu à Justine, Quentin et Hélène pour leur accueil lors de mes tous premiers jours au labo. Merci à toute la *clique* du labo : Olivier, Irving, Nina, Bin, Clémentine, Lily, Laurence et Sharon. Merci aussi à Alexandre d'avoir boosté ma thèse en seulement 3 mois de stage ! Merci aussi à Maxim et Mathilde sans qui aucune des expés de la thèse n'aurait été possible.

Merci Sandrine de m'avoir accueilli et accompagné, tant bien que mal, dans cette aventure. Ton énergie positive m'a aidé à ne rien lâcher dans les moments les plus difficiles ! Merci à Franck de m'avoir permis de lancer ce projet. Un grand merci aux doctorants (Elsa, Arthur, Caliani et Alexandra) de m'avoir rendu cette période un peu moins désagréable. Merci aussi à toute l'équipe Sciences (en particulier Joan, Valentine et Mickaël), à l'équipe lyonnaise (en particulier Laurent, Leslie, Laurène et Olivier) et aux *humans* (en particulier Cédric et Amandine) de m'avoir fait une place parmi vous.

Merci à tous mes amis de m'avoir permis de sortir la tête de l'eau et de me rappeler à quel point la thèse, ce n'est pas la vraie vie. Chaque moment passé avec vous m'a aidé à penser à autre chose et à me remettre au travail avec un peu plus de motivation. Même si on s'est peu vus ces dernières années, sachez que vous voir plus souvent était une de mes sources de motivation pour finir cette thèse. Merci à Hugo, Matthieu, Matthew, Ronan & Justine, Sandrine & Léa, Louis, lauP & jenB, Stéphanie, Apolline, Maëlys & Vincent et Lyamani.

Merci à toute à ma famille et en particulier mes parents, Nicole & Nestor et mes sœurs, Alice & Auristelle. Je sais que cela n'a pas toujours été facile de comprendre ce que je faisais pendant mes études. Je ne peux pas vous promettre que ce sera plus facile de comprendre ce que je vais faire maintenant, mais je sais que je pourrai toujours compter sur votre soutien et votre amour. J'espère faire au mieux avec le meilleur de ce que vous m'avez transmis. Malgré ma pudeur et ma façon de communiquer parfois insuffisante, je tiens à dire ici que je suis heureux de faire partie de notre famille.

Marie, je pense te devoir toutes mes réussites, car tu m'offres une confiance en moi inépuisable, simplement en posant ton regard sur moi. Je n'aurai jamais osé me prétendre capable de mener à bien un doctorat sans te savoir à mes côtés. Merci pour ta patience, malgré toutes ces soirées, ces week-ends et ces vacances sacrifiées pour la thèse. Merci de m'avoir fait une place dans ta famille, qui m'a si bien accueilli. Merci de m'avoir soutenu, protégé, consolé, conforté, épargné et aimé malgré cet envahissant projet de presque 4 ans. Merci d'être là pour moi, chaque jour, depuis 6 ans.

Merci (aussi) au moi d'il y a 9 ans d'avoir osé me réorienter en sciences cognitives. Merci au moi d'il y a maintenant 4 ans d'avoir parié que je pouvais convaincre un laboratoire et une entreprise de me suivre dans un tel projet. Merci au moi de ces 3 dernières années d'avoir maintenu le cap ! 😊

## Acknowledgements

First of all, thank you, Jean-Christophe and Vincent, for giving me this opportunity and trusting me. Among PhD students, it is often said that the supervision plays a crucial role, not only in guiding the research but also in maintaining the mental health of the doctoral candidate. I believe I was able to successfully complete my thesis while staying in a good mental space thanks to you. I hope everyone pursuing a PhD gets to be supervised as I was. I am grateful for the trust and the relationship we built, which extended far beyond the scope of the thesis. Thank you very much!

Thank you to Angela Sutan, Alexandre Zénon, Wim de Neys, and Antoine Mandel for agreeing to be part of my jury. I hope that reading my thesis over the holiday season might serve as an excuse to escape any awkward situations... It's the only advantage I could find; sorry for the inconvenient timing! Thank you also to Vincent Berthet, Jérôme Sackur, Isabelle Régner, Sandrine Bélier and Thomas Breda for being part of my thesis committees. Your advice and support for my work have greatly helped me in completing this thesis. A special thanks to Vincent, who inspired me and sparked my interest in cognitive biases and bounded rationality back when I was still an undergraduate student in Nancy.

Thank you to everyone who helped shape my professional path up to the completion of this thesis. Special thanks to Tania and Mohamed from L'École du Recrutement, who helped me find recruiters for my experiments; Julie Grèzes and Rocco Mennella, who gave me my first internship in a lab; Cog'X, which allowed me to explore cognitive science in a professional setting; and Aurélien & Hanaë, who gave me the chance to present my thesis during a TEDx event.

Thank you to my fellow PhD students, who made me feel like I belonged. Nothing against people who chose different paths in life, but there is something unique about the bond we share. Initially, I thought coming to the lab was only useful for meetings with my supervisors. What a mistake! Over time, I realised how fortunate I was to be surrounded by you, to get to know you, to rely on your support, and to help each other. Special thanks to Justine, Quentin, and Hélène for welcoming me during my very first days in the lab. Thank you to the entire lab crew: Olivier, Irving, Nina, Bin, Clémentine, Lily, Laurence, and Sharon. And thank you to Alexandre for giving my thesis a boost in just three months of internship! Thanks also to Maxim and Mathilde, without whom none of the thesis experiments would have been possible.

Thank you, Sandrine, for welcoming and supporting me throughout this journey, through its ups and downs. Your positive energy helped me stay resilient during the toughest times! Thanks to Franck for enabling me to embark on this project. A big thank you to the PhD students (Elsa, Arthur, Caliani, and Alexandra) for making this period a little less unpleasant. Thanks also to the Science team (especially Joan, Valentine, and Mickaël), the Lyon team (especially Laurent, Leslie, Laurène, and Olivier), and the Humans team (especially Cédric and Amandine) for making room for me among you.

Thank you to all my friends for helping me keep my head above water and reminding me that the PhD is not real life. Every moment spent with you helped me think about something else and return to work with a bit more motivation. Even though we didn't see much of each other in recent years, know that spending more time with you was one of my sources of motivation to finish this thesis. Thank you, Hugo, Matthieu, Matthew, Ronan & Justine, Sandrine & Léa, Louis, lauP & jenB, Stéphanie, Apolline, Maëlys & Vincent, and Lyamani.

Thank you to all my family and in particular my parents, Nicole & Nestor and my sisters, Alice & Auristelle. I know it hasn't always been easy to understand what I was doing during my studies. I can't promise it will be easier to understand what I'll do next, but I know I can always count on your support and love. I hope to do my best with the values you've passed on to me. Despite my reserved nature and my sometimes-inadequate communication, I want to say here that I am proud to be part of our family.

Marie, I owe all my achievements to you because you provide me with unwavering confidence simply by looking at me. I would never have dared to believe I was capable of completing a PhD without knowing you were by my side. Thank you for your patience, despite all those evenings, weekends, and holidays sacrificed for the thesis. Thank you for welcoming me into your family, who received me so warmly. Thank you for supporting, protecting, consoling, comforting, sparing, and loving me despite this all-consuming project of nearly four years. Thank you for being there for me every day for the past six years.

Thank you (also) to my former self from nine years ago for daring to switch to cognitive science. Thank you to my self from four years ago for betting that I could convince a lab and a company to support such a project. And thank you to my self over these past three years for staying the course! 😊

# Table of contents

<b>General Introduction.....</b>	<b>1</b>
<b>Chapter 1: Exploration Bias Towards Favourites in Set Selection .....</b>	<b>32</b>
<b>Chapter 2: Cognitive Factors Influencing Information Search Biases .....</b>	<b>71</b>
<b>Chapter 3: Metacognitive Judgements in Exploration Strategies for Set Selection .....</b>	<b>104</b>
<b>Chapter 4: Learning and Adapting Optimal Strategies .....</b>	<b>122</b>
<b>General discussion .....</b>	<b>187</b>
<b>Résumé substantiel en français .....</b>	<b>218</b>

# General introduction

*“In an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.”*

— Herbert A. Simon

*‘Designing Organizations for an Information-Rich World’ in Martin Greenberger (ed.) Computers, Communications, and the Public Interest (1971)*

Approximately 252,000 new websites are created daily (Drezner & Edigbe, 2024). According to Statista, 347 billion emails are sent globally each day. One person receives 43 emails per day on average (Statista, 2024). As Herbert Simon noted, the ability to allocate attention efficiently among a plethora of information is critical. In this thesis, we focus on how to optimise information search for making a complex decision that involves selecting multiple options simultaneously rather than just one.

The introduction is organised into six sections. In the first section, we present the origins of the project that underpin the work conducted throughout the thesis. The second section highlights the critical differences between single-option and set choices, leading to our research questions. In the third section, we introduce the original experimental paradigm employed in this thesis. The fourth section reviews the existing paradigms in literature on information search and decision-making. The fifth section review some relevant biases and cognitive limitations. Finally, the sixth section outlines the structure of the thesis chapters.

## Table of contents

Table of contents .....	2
I. Origin of the Thesis Project: Decision-Making Among Recruiters.....	4
a. Sourcing: talent hunting on LinkedIn -----	4
b. Previously, in recruitment studies... -----	5
c. Adjusting Research Focus -----	6
II. Single Selection vs. Set Selection: What Difference Does It Make? .....	6
a. Choose a Single Option vs. Choose a Set of Options -----	6
b. The Need for Adapted Information Search Strategy -----	7
c. The Risk of Biases and Heuristics in a Complex Problem -----	8
d. Research Question -----	9
III. Development of an Experimental Paradigm for the Study of Information Search in Set Selection.....	9
a. Set Exploration – Set Selection Paradigm -----	9
b. Key Features of the Paradigm -----	10
IV. Information Search and Decision-Making: A Review.....	12
a. Secretary Problem paradigm -----	13
b. Exploration – Exploitation dilemma -----	14
c. Sequential evidence accumulation paradigm -----	16
d. Multi-attributes decision-making paradigm (MADM) -----	18
V. Various sources of bias and mistakes.....	20
a. Confirmation biases -----	20
b. Pavlovian Biases -----	21
c. Uncertainty Aversion -----	22
d. Choice Bracketing -----	23
e. Failure of Contingent Thinking (FCT) -----	24
VI. Structure of the Thesis.....	25
a. Chapter 1: Exploration Bias Towards Favourites in Set Selection -----	25
b. Chapter 2: Cognitive Factors Influencing Information Search Biases -----	25

General introduction

c.	Chapter 3: Metacognitive Judgements in Exploration Strategies for Set Selection	26
d.	Chapter 4: Learning and Adapting Optimal Strategies -----	26
	Bibliography.....	27

# I. Origin of the Thesis Project: Decision-Making Among Recruiters

To begin, we describe the starting framework of the thesis, which focuses on recruiters' decision-making processes. We explain how we started from a recruitment issue, known as sourcing, to study set selection in a controlled laboratory environment through this thesis.

## a. Sourcing: talent hunting on LinkedIn

Sourcing is a recruitment practice that involves reaching out to candidates rather than waiting for responses to a posted job offer. This practice is part of the recruitment techniques that are developing, particularly thanks to social networks like LinkedIn, which allow recruiters to approach passive candidates, that is, individuals currently employed but open to new opportunities (Hosain & Liu, 2020). According to a study conducted by *LinkedIn* in 2015, 75% of workers considered themselves passive candidates.

One of the first challenges for recruiters is being able to filter relevant profiles before establishing contact. LinkedIn claims to have over 1 billion users worldwide in 2024, including 30 million in France. In just two years, the number of French users is said to have increased by 20%. A simple search for "Data Engineer" based in France yields 46,000 profiles. Reviewing all these profiles is clearly impossible, and even going through the details of just the top 10 profiles displayed on the first results page is too time-consuming.

Another challenge is identifying the relevant information to explore. Each profile contains multiple sections (experience, education, activities on the site, etc.), which can each be highly detailed. Every user completes their own profile with varying levels of detail. A recruiter must sift through all the available information to find the relevant details to determine whether a profile is of interest for the recruitment process. In general, it is impossible for a recruiter to explore all the relevant information for every profile that might be of interest. The goal is to identify the challenges recruiters may face in this practice, which is becoming an increasingly significant part of their role. According to *LinkedIn's* 2015 study, 61% of companies hired passive candidates, which requires efficient exploration to create candidate shortlists.

Studies show that sourcing on the web, particularly on LinkedIn, is rapidly expanding (Abbas et al., 2021). This raises many questions about how candidates are screened and selected from the vast array of relevant and accessible profiles online (D'Silva, 2020; Roulin & Levashina, 2019).

## b. Previously, in recruitment studies...

To our knowledge, the scientific literature does not yet tell us whether recruiters' decision may suffer from suboptimalities due to cognitive skills. A few studies have pointed to suboptimal CV reading processes (Cole et al., 2007), but most recruitment research focuses more on the social aspect (such as gender gap or racial discrimination) rather than the cognitive aspect. Studies in economics, psychology, and sociology on recruitment have clearly demonstrated systematic discrimination typically caused by stereotypes, a form of bias. Audit and correspondence studies<sup>1</sup> have quantified discrimination in the labour market, particularly showing how characteristics like names or ethnicity influence the likelihood of being called back for an interview. Numerous findings unanimously indicate that minorities are systematically disadvantaged in the labour market (Eyting, 2022; Jowell & Prescott-Clarke, 1970; Koch et al., 2015; Kroll et al., 2021; Lang & Lehmann, 2012; Neumark, 2018; Pager et al., 2009). Future studies should focus more on exploring solutions to reduce these discriminations in the labour market (Bertrand & Duflo, 2017).

A 2016 study provided an in-depth theoretical and experimental analysis of the concept of attentional discrimination, which posits that decision-makers, such as recruiters, allocate their attention endogenously and in a biased manner depending on stereotypes associated with an ethnic or social group and the type of market (Bartoš et al., 2016). The article demonstrates that candidates subject to stereotypes receive less attention in so-called “cherry-picking” markets—situations where significant filtering occurs between two stages. For instance, in the labour market, significant filtering happens between the application phase and the interview stage. Consequently, candidates subject to stereotypes, whose applications receive less attention and exploration, have lower chances of being among the few invited to an interview. Conversely, stereotyped individuals receive more attention in “lemon-dropping” markets—situations where there is little filtering between one stage and the next. For example, in the housing rental market, few applicants are rejected between initial contact and viewing appointments. Here, stereotyped candidates face a higher likelihood of being among the few prevented from scheduling a viewing because their profiles receive more attention in these situations.

These results highlight that discrimination can manifest as early as the information search phase, where recruiters apply filtering rules under the influence of stereotypes. This biased

---

<sup>1</sup> Audit and correspondence studies, methodologies where fictitious applications with varying characteristics (e.g., names, gender, age, address) are sent to employers to measure differential treatment.

attentional discrimination based on market type conditions information search during set selection, thereby exacerbating inequalities in everyday situations.

### c. Adjusting Research Focus

Our initial project aimed to study real recruitment processes (evaluate recruiters' cognitive biases and their efficiency in using assessment IT tools) within French companies, through a partnership with a consulting firm. However, all our proposals were declined for various reasons. It turns out that recruitment processes are less standardised than one might expect. Among the companies we approached, very few had archives of past applications or recruitment procedures, making it impossible to reconstruct such processes. Other companies declined due to concerns about granting access to sensitive internal data. As a result, we decided to focus on laboratory experiments only.

While the study of social biases like stereotypes is well-established in the literature (Bordalo et al., 2016; Cauthen et al., 1971; Ellemers, 2018; Sue & Kitano, 1973), there remain significant unknowns about how recruiters handle the overwhelming amount of information available for sourcing. Expertise in information search aimed at creating a shortlist appears to be a particularly relevant skill for recruitment professionals. The starting point of our paradigm is the following question: how should information about candidates be sought when it is impossible to collect all the information, and a shortlist must be created? We designed an experimental framework to create an original recruitment task that captures the complexity of a practice serving as the starting point for many recruitment processes.

## II. Single Selection vs. Set Selection: What Difference Does It Make?

Before presenting our paradigm in detail in the next section, we focus here on the fundamental differences between making a single choice and making a set selection. As we shall see, set selection is a complex problem where search strategy might be counterintuitive, and where the risk of errors and biases is heightened compared to single selection.

### a. Choose a Single Option vs. Choose a Set of Options

Decision-making is an ubiquitous process in daily and professional life, often involving information search before making a choice. This could be a student selecting the courses they want to take, a vacationer deciding on activities, someone choosing ice cream flavours, or a

school choosing its future students. The common thread among these examples is that they describe situations where individuals select a set of options. These situations differ from others, such as a person choosing a new phone, a family deciding on their next holiday destination, or someone selecting a movie to watch.

Set selection involves multiple potential errors that do not arise in single-choice selection. Forming a team of 10 requires considering not only each individual element but also the interdependencies between them, as well as the need for cohesion or complementarity. For instance, a student selecting courses for their curriculum must ensure that the chosen courses are compatible in terms of scheduling. In team sports, a coach must build their team while paying attention to the cohesion between players. Similarly, choosing a set of ice cream flavours is not just about selecting one's favourites but also about ensuring that the chosen flavours complement each other.

One might assume that the consequences of a poor choice are less severe in set selection compared to single selection. Unfortunately, this is not always the case. Making one bad choice out of 10 might seem less critical than making one bad choice out of one, as the remaining 9 good elements could theoretically compensate for the mistake. However, if the good elements cannot offset the errors, the entire selection can fail. A team with members who undermine its cohesion can lose a match, a single bad flavour can ruin the entire ice cream combination, and overlapping courses in a schedule can disrupt all of your planning. Ultimately, the risk of failure is greater in set selection than in single selection.

### b. The Need for Adapted Information Search Strategy

To make the best choice, the preliminary information search is crucial and differs between single and set selection. For set selection, it is not merely a matter of identifying the single best option. Suppose a student needs to choose five courses from ten available options to validate their year. They may consult the syllabus for each course and attend course presentations. However, they cannot attend all presentations and must rely solely on the syllabus for five of the ten courses. Which course presentations should the student attend? If they had to choose a single course, the answer is straightforward: after consulting the syllabi, they should only attend the presentations of the most appealing courses. This approach is intuitive and optimal.

But what if they need to choose five courses? Intuitively, one might employ the same strategy—attending the presentations of the five most appealing courses. However, this

strategy, while intuitive, is no longer optimal. In general, the information search must enable the identification of the best set of options. In a context constrained by time, financial resources, or computational capacity, the information search should focus on the most decisive information for set selection. In the case of selecting five courses from ten, the student would benefit most from attending the presentations of courses they are most uncertain about. Attending presentations of courses they are already inclined to choose or almost certain to reject offers little value in finalising their selection. In other words, in the context of set selection, focusing on the most preferred option becomes suboptimal which is largely counter-intuitive, and markedly different from the context of a single selection. Information search strategies can therefore differ profoundly between single and set selection.

### c. The Risk of Biases and Heuristics in a Complex Problem

As the decision-making and information search literature shows, individuals exhibit bounded rationality (Simon, 1955), characterised by cognitive biases and the use of heuristics. Cognitive biases are generally described as systematic and intuitive errors in reasoning compared to a logical or mathematical solution. For example, participants are told that there is a group of 100 people, 30 of whom are engineers. When a person (named Jack) is described as someone who enjoys puzzles, participants tend to overestimate the probability that Jack is one of the engineers. However, their estimates are accurate when no description of Jack is provided. The associated heuristic is the "representativeness heuristic", which involves judging the likelihood of an event belonging to a category based on its similarity to the prototype of that category. This leads to a cognitive bias known as "base rate neglect", where individuals tend to ignore prior probabilities when evaluating the likelihood of an event (Kahneman & Tversky, 1973). From another perspective, heuristics are strategies that allow individuals to make effective use of their cognitive resources by simplifying the problem before solving it (Gigerenzer et al., 1991). We will return to biases and heuristics in sections IV and V.

The size of the selection also increases computational complexity. If the complexity of a single choice often exceeds our cognitive capacities, it becomes even more challenging in set selection. While choosing 1 option from 10 requires evaluating and comparing 10 expected outcomes, selecting 2 options from 10 requires considering 45 possible pairs. For a student choosing 5 courses from 10, there are 252 possible combinations. The challenges posed by set selection have been underexplored, even though it is clear that our bounded rationality may prevent us from solving such problems perfectly.

#### d. Research Question

In sum, set selection differs from single selection because the optimal search strategy required is fundamentally different and counterintuitive. Set selection can involve greater risks or necessitate accounting for interdependencies. Finally, set selection is computationally more complex because it requires considering a significantly larger number of possibilities, and might be more prone to suboptimal or biased strategies because of this complexity.

The literature has largely overlooked how set selection challenges our cognitive capacities. This thesis aims to fill this gap and to understand the strategies individuals use to make set selections. Specifically, we seek to understand the exploration strategies employed to choose which information to examine and how individuals make set selections despite uncertainty regarding some of the options.

### III. Development of an Experimental Paradigm for the Study of Information Search in Set Selection

In this section, we describe the experimental paradigm we developed and used throughout the thesis. Several versions of the task were employed. Here, we present the baseline version that served as the foundation for the thesis, detailing its characteristics and the research framework within which it was applied. The derived versions of the task are introduced in the respective chapters.

#### a. Set Exploration – Set Selection Paradigm

The paradigm involves two stages: an exploration phase followed by a selection phase. After participants are provided with initial information (an integer score A between 0 and 10) for each of the 10 available options, they proceed to the exploration phase. In this exploration phase, they select 5 options for which they will receive additional information (also an integer score B between 0 and 10). After obtaining the second score for these 5 options, participants then enter the selection phase. At this stage, they select the 5 options out of the initial 10 that they believe to be the best (based on the sum of scores A and B).

Stemming from the initial idea of studying recruitment decisions, the options in the paradigm represent profiles of individuals who are potential candidates for an unspecified position. Scores A and B represent distinct skills and are deliberately uncorrelated (which

participants are told). The final selection at the second stage corresponds to the set of recruited candidates.

### b. Key Features of the Paradigm

To describe the key features of our paradigm and to compare it with other paradigms in the literature in the next sections (Section IV a–d) we will rely here on a classification proposed recently by (Gigerenzer, 2024), presented in Table 1. It aims to distinguish experimental paradigms using concepts such as risk, ambiguity, uncertainty, and intractability, to avoid confusion between them. Based on terms used by Savage (1954, 1972), this classification separates “Small worlds,” where an optimal solution is calculable, from “Large worlds,” where an optimal solution is either too costly or impossible to calculate. In a risk task, all consequences and probabilities associated with future states are known, making it possible to calculate an optimal action (e.g., lottery games). In an ambiguity task, the future states are known, but the probabilities associated with those states are not (e.g., biased dice with unknown probabilities). Both terms fall under “Small worlds”. Tasks involving uncertainty feature unknown future states and probabilities (e.g., the introduction of an innovative product to the market). Finally, tasks involving intractability have known states and probabilities, but the computation of an optimal solution is impractical due to the problem’s complexity (e.g., finding the shortest route in a Traveling Salesman Problem with a large number of cities<sup>2</sup>).

Conditions	Small worlds		Large worlds	
	Risk	Ambiguity	Uncertainty	Intractability
Are all possible future states and consequences of all actions known?	Yes	Yes	No	Yes
Are all probabilities known?	Yes	No	No	Yes
Can optimal action be calculated?	Yes	Yes	No	No

Table 1. Risk, ambiguity, uncertainty, and intractability (Gigerenzer, 2024)

Our paradigm fits into a "Large world" type, characterised by intractability—a context where the optimal strategy cannot be computed, even if it theoretically exists (see Table 1). In detail, our paradigm comprises four essential features: simultaneous actions, computational

---

<sup>2</sup> The Traveling Salesman Problem (TSP) is a classic combinatorial optimisation problem in mathematics and computer science. It involves finding the shortest route for a traveller to visit a given list of cities exactly once each and return to the starting city. This problem is well known for its computational complexity, as the number of possible solutions grows exponentially with the number of cities. Currently, there is no known algorithm capable of solving it optimally in polynomial time for a sufficiently large number of cities (Hoffman & Padberg, 2001).

complexity, the presence of an intuitive yet suboptimal heuristic, and a non-intuitive but nearly optimal heuristic. Each feature is discussed below.

### **Simultaneous Actions**

First, the paradigm captures exploration and set selection through simultaneous actions at each stage. Participants must always compose a set before advancing to the next stage. This constraint requires participants to determine a set before knowing the outcomes of their actions.

### **Limited Information Access**

Second, access to information is strictly limited, as is the size of the selection set. Participants must explore and select exactly the number of options specified (5 explorations and 5 selections in the baseline version of the task). They never have access to 100% of the information, but they also cannot decide to acquire less information. Consequently, information acquisition is cost-free, there are no time constraints for completing the stages, and the selection phase always includes unknown B scores. However, participants are free to choose for which options they wish to obtain information.

### **Computational Complexity**

Third, the paradigm is computationally complex. Calculating the expected gains for each exploration and selection possibility to determine the optimal strategy is computationally prohibitive, even for standard laboratory computers. For instance, with 5 options to be chosen out of 10, there are 252 possible set combinations, compared to only 10 combinations in a single-choice scenario. If participants must explore 5 options out of 10 before finding the best set of 5, there are 252 exploration combinations, each followed by 252 set choices, resulting in 63,504 potential pathways. Furthermore, each option comprises two scores with 11 possible values (integers ranging from 0 to 10). The computational complexity is discussed in greater detail in Chapter 1.

### **Intuitive and Nearly Optimal Heuristics**

Fourth, the paradigm allows participants to employ heuristics. Specifically, there is an intuitive heuristic that is clearly suboptimal. This heuristic involves exploring the options with the highest A scores, as mentioned previously in the example of a student choosing their courses. There is also a non-intuitive but nearly optimal heuristic, which involves exploring the options near the middle ranks (based on scores A), as they are most likely to shift into or

out of the final selection set. For instance, in the case of a student selecting 5 courses, the best heuristic would involve attending the presentations for courses ranked near the 5th position after reviewing all course descriptions. In the experiment, it is thus more effective to explore the scores B of candidates close to the 5th rank according to their scores A. This heuristic can be adjusted to accommodate the size of both the exploration and selection sets (see Chapters 1 and 4 for different task versions). These heuristics can be applied on a trial-by-trial basis without requiring learning across trials, as all problem data are known.

### **Research Framework of the Paradigm**

As this paradigm is novel and addresses an original question in the literature, the thesis focuses on a few fundamental research questions related to set selection. We prioritise examining the gap between the optimal solution and participants' strategies. We describe participants' behaviours using simple heuristics to draw initial conclusions. Observed phenomena are compared to several cognitive biases well-documented in the literature. Finally, we aim to understand how individuals perceive their own strategy for solving the paradigm before testing potential remediation methods.

As detailed further in the General Discussion of this thesis, these preliminary studies constitute a first step, which could later enable the investigation of related questions, such as the non-instrumental value of information or the psychological determinants of inter-individual differences (Kelly & Sharot, 2021; Kobayashi et al., 2019; Kobayashi & Hsu, 2019; Sharot & Sunstein, 2020).

## **IV. Information Search and Decision-Making: A Review**

There is a vast experimental and theoretical literature in cognitive psychology on information search in decision-making. This literature has primarily focused on single-choice scenarios rather than set selection. Here, we review four paradigms that have been particularly prominent in this field and highlight how they differ from our approach. For each paradigm, we provide its origin, definition, a categorisation of the type of "world" studied according to the framework proposed by (Gigerenzer, 2024), and the insights the paradigm offers on individuals' decision-making processes.

## a. Secretary Problem paradigm

### Presentation

The Secretary Problem is a decision-making paradigm with sequential information acquisition, first introduced by statistician Lindley D.V, in 1961 (Ferguson, 1989; Freeman, 1983). The central question is how to choose the best option when options are presented one at a time, and rejecting one is necessary to access the next. In this task, a decision-maker observes a sequence of candidates for a secretary position and must choose the best without the ability to revisit previous options. The candidates are presented in random order, and the decision-maker learns a candidate's value only when they are presented with that candidate. The decision-maker must then either choose the candidate, thereby ending the process, or reject the candidate permanently to consider the next candidate. This process can continue as long as candidates remain.

### Classification

According to Gigerenzer, (2024) categorisation, the Secretary Problem represents a small ambiguous world. This categorisation is established by addressing three questions outlined by Gigerenzer. *Are all possible future states and consequences of all actions known?* In most studies, the consequences of the two possible actions are known: stopping the search provides a reward corresponding to the chosen option, while continuing allows access to a new option with a predetermined cost. The value of the new option (and possibly future ones) is not known in advance. *Are all probabilities known?* The exact probabilities of each option's occurrence are unknown, which constitutes the primary difficulty of the task. *Can optimal action be calculated?* It is possible to compute an optimal stopping strategy, as demonstrated by mathematical studies. This strategy involves rejecting slightly more than the first third of candidates and then selecting the first candidate superior to all previously evaluated candidates. This approach yields the best candidate with a probability of  $p = .378$  (Chow et al., 1964). Subsequent work has shown that optimal stopping rules can also be calculated under various conditions, such as when the number of candidates is uncertain (Gianini-Pettit, 1979; Rasmussen & Robbins, 1975; Stewart, 1981), when the value range is unknown (Mucci, 1973), when the candidate can turn off the offer (Petrucci, 1981; Smith, 1975) or when previous candidates may still be selectable (Corbin, 1980; Petrucci, 1981). The literature in applied mathematics and statistics has demonstrated that optimal stopping rules can be established in diverse scenarios. It was mostly used for representing decisions such as recruitment, negotiation, or marriage.

In behavioural economics and experimental psychology, the Secretary Problem has revealed that individuals can use heuristics to solve the task satisfactorily (Corbin, 1980), although their strategies consistently deviate from the optimal. Participants' strategies are often described by heuristics. The Cut-off Rule describes a strategy close to the optimal one: rejecting a set number ( $r$ ) of candidates and then choosing the first candidate who surpasses all previously evaluated ones. Applying this rule with  $r$  set to the optimal level achieves the best possible outcome. However, individuals tend to under-sample (Seale & Rapoport, 1997), especially when the cost of accessing new candidates increases (Rapoport & Tversky, 1970). This under-sampling, where participants' cut-off tends to be smaller than the optimal cut-off, can be viewed either as a bias—reflecting systematic errors that prevent optimal performance—or as an ecological heuristic, maximising performance given the cognitive effort required for memorising information and calculations.

### **Comparison with the set-selection paradigm**

The Secretary Problem could be adapted to study information search preceding set selection, which, to our knowledge, has not yet been done in behavioural economics or cognitive psychology. An optimal strategy could be computed if it were possible to make multiple choices during the sampling process (Kleinberg, 2005). In this version, options would still be sampled sequentially, and selecting one option would not end the process, which would continue until the participant had chosen  $N$  options. However, exploration would remain completely tied to selection, as rejecting one option would still be necessary to acquire more information. While this characteristic would be less prominent in multi-selection contexts, it would still apply to the final selection. This is the main difference with our paradigm.

## **b. Exploration – Exploitation dilemma**

### **Presentation**

The exploration–exploitation paradigm addresses the dilemma faced by individuals or systems when deciding whether to explore new options to acquire additional information, potentially improving future outcomes, or to exploit known options to maximise immediate rewards (Bellman, 1956; Bellman & Kalaba, 1957).

A common task for studying this dilemma is the Multi-Armed Bandit Task, where participants face several options, often likened to slot machines (Fig. 1). Each option, or "arm," has a different probability of yielding a reward. Participants must select an arm on each trial, attempting to balance exploration—testing new arms to learn about their

probabilities of reward—and exploitation—choosing the arm with the best known return to maximise immediate gains (Kaelbling et al., 1996). For example, one machine might offer a reward 80% of the time, another 50%, and another 30%, but these probabilities are initially unknown. The task is to determine which machines are most profitable while accumulating rewards.

## Classification

This task represents a large uncertain world under Gigerenzer (2024) categorisation. The consequences of future actions are unknown, probabilities are not provided, and an optimal solution cannot be calculated. In simplified versions of the task, where reward probabilities are stable for each arm, the optimal action at each turn can be determined mathematically by calculating the Gittins index (Gittins, 1975, 1979; Gittins & Jones, 1979; Weber, 1992). However, this simplified version does not reflect real-world scenarios, such as selecting research topics or deciding on company projects. In more complex versions of the task, where reward probabilities evolve over time, and more than two arms are available, no calculable optimal strategy has been identified (Cohen et al., 2007; Daw et al., 2006). The computational demands of solving even the simplified versions far exceed human capacity. This distinction makes the Multi-Armed Bandit Task significantly different from the paradigm developed in this thesis.

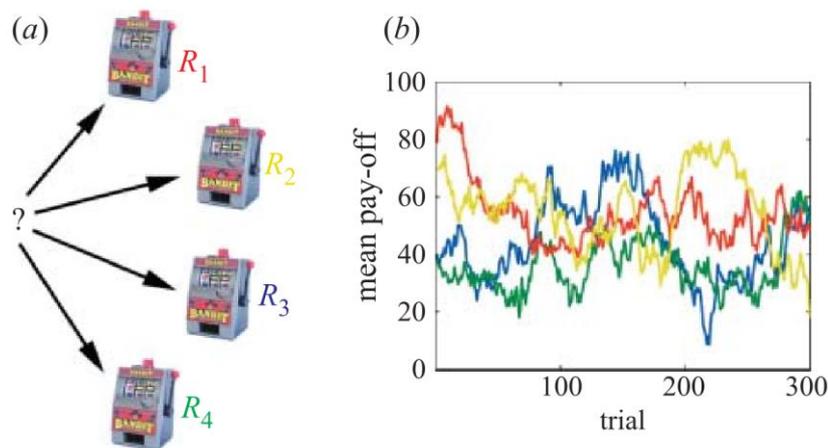


Fig. 1 Illustration by Cohen et al. (2007) depicting the task designed in the study by Daw et al. (2006). Panel (a) displays the four available slot machines for each trial of the task. Panel (b) shows the evolution of each arm's value over the course of the trials.

The study of strategies employed by individuals is describes by various heuristics that have been proposed. Individuals might predominantly exploit the best option while occasionally exploring randomly—a strategy known as the epsilon-greedy heuristic (Sutton & Barto, 1998). Alternatively, they might exploit the best options but with noise —known as the softmax heuristic (Thrun, 1992). According to a review comparing these hypotheses,

individuals' behaviour is better described by the softmax heuristic (Cohen et al., 2007). Inter-individual variations may be influenced by psychiatric factors. For example, a psychiatric disorder linked to addiction could correspond to an increased tendency towards exploitation, whereas a depressive disorder might bias individuals towards maladaptive exploration (Addicott et al., 2017).

### **Comparison with the set-selection paradigm**

As in the Secretary Problem paradigm, it is not truly possible to distinguish information search from decision-making within the exploration–exploitation paradigm as the decisions made in each trial contribute to acquiring information for subsequent trials. This is one of the key differences with our paradigm. Also, in our task, participants start each trial with the same information, and trials are independent of one another.

One could imagine a version of the task where participants are allowed to choose multiple arms in each trial, making the task more ecologically valid. We will return to this research possibility in the General Discussion.

### **c. Sequential evidence accumulation paradigm**

#### **Presentation**

Sequential evidence accumulation paradigms involve receiving information incrementally and stopping the process when a decision is ready to be made. The decision typically involves choosing an option from a set. For example, one task might involve determining which horse will win a race among four horses (A, B, C, or D), as illustrated in Fig. 2. The participant receives information 1, which is 75% reliable, indicating horse A. The participant can then choose to pay (in points) to acquire information 2 or they can stop accumulating information and make their choice of which horse to bet on (Hausmann & Läge, 2008).

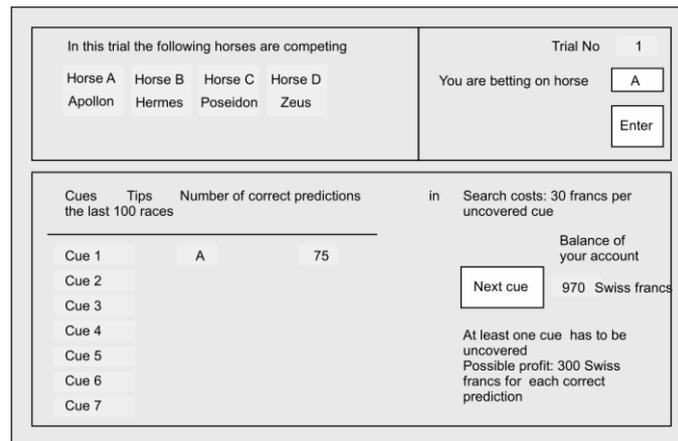


Fig. 2. Image extracted from Hausmann & Läge (2008). Representation of a trial in an evidence accumulation task. The first cue is displayed (middle-left section of the screen) indicating "bet on horse A" with a reliability of 75%. The top-right section of the screen shows the participant deciding to bet on horse A without accumulating additional evidence. The lower-right section indicates that the cost of accessing further information is 30 Francs, the participant possesses 970 Francs, and they can win 300 Francs if their bet on the correct horse is successful.

## Classification

This type of task is classified as a small uncertain world according to Gigerenzer's (2024) categorisation. All possible actions are known: sampling provides additional information at a fixed cost, and making a decision results in a known reward for success and a known loss for failure. However, the exact probabilities of success are unknown because the probability of unrevealed cue are unknown. Furthermore, it is possible to determine a rational model for this task. To do so, one must calculate the expected gain based on the known information and compare it to the expected gain of acquiring an additional piece of information, considering the expected value of that new information (Lee & Cummins, 2004).

Heuristics are widely used to describe individuals' behaviours when faced with this paradigm. The primary objective of this field of research is to understand how individuals define their stopping rules for evidence accumulation. As described by Hausmann & Läge (2008), part of the literature suggests that individuals rely on single pieces of information (One-Reason Decision-Making, ORDM) rather than integrating multiple pieces (Multiple-Reason Decision-Making, MRDM). The most well-known ORDM heuristic is Take-The-Best (TBB), which involves relying on the information associated with the best probability (Gigerenzer et al., 1991). Rational models, by contrast, are generally categorised under MRDM, as they usually require accumulating at least two pieces of evidence to compute the stopping rule. In cases where it is optimal to decide after only one piece of information, rational models may actually be outperformed. Some studies propose unifying these heuristics with a stopping rule based on a desired level of confidence (DLC). This criterion would not be fixed by the number of cues or a general rule but by a personal confidence threshold that determines when

to stop accumulating evidence (Hausmann & Läge, 2008). This theory allows researchers to explore the significant inter-individual variability observed in previous studies (Lee & Cummins, 2004; Newell et al., 2004).

### **Comparison with the set-selection paradigm**

An important difference with our paradigm (in addition to the single selection aspect) is the possibility of directing the search for information towards a desired option. In this paradigm, the participant does not know which option he will have information on when he chooses to accumulate additional information (the next information could be about horse A or it could be about horse B, C or D). In our paradigm, the participant decides to explore a particular option, and gets information about that option. With certain adaptations, this paradigm could be employed to study information search preceding set selection, as we will see in the general discussion.

#### **d. Multi-attributes decision-making paradigm (MADM)**

### **Presentation**

The final paradigm we present is the Multi-Attribute Decision-Making (MADM) paradigm. This framework is used to understand decisions involving multiple dimensions or attributes, each with a specific importance or validity. Typically, participants are faced with several alternatives described by attributes with varying predictive weights. For example, in the work of Söllner et al. (2013), a typical task involves evaluating options where cues A, B, C, and D have distinct predictive weights (Fig. 3). Based on this information, participants must choose the best option. Unlike the sequential evidence accumulation paradigm, where information is gradually revealed and may involve search costs, MADM tasks present all information simultaneously at the start of each trial. In this paradigm, the extrinsic costs associated with gathering cues are also eliminated, as highlighted by Söllner et al. (2014). In this paradigm, the search for information is therefore endogenous and mainly visual, which can be investigated using eye-tracking (Meißner et al., 2020).

**Which city has more inhabitants?**  
 Validities: A: 80%, B: 70%, C: 60%, D: 55%

	Garango	Bingo
A)	-	+
B)	+	-
C)	+	-
D)	+	-
	Choose	Choose

Fig. 3. Example of a trial from a multi-attribute decision-making (MADM) task used by Söllner et al. (2013). The cues A, B, C, and D have varying predictive validities regarding the correct answer. An optimal participant is expected to calculate exactly which option is most likely to be correct based on the cues' directions (+ and -) and their validities.

### Classification

According to Gigerenzer (2024) categorisation, the MADM paradigm can be classified as a small risky world, where all alternatives and their consequences are known. The possible actions are typically limited to selecting one of the available options, with associated gains or losses explicitly defined beforehand. In these scenarios, an optimal solution is formally identifiable: it involves calculating the weighted sum of the cues for each option and selecting the one with the highest value.

Current theories on multi-attribute decision-making revolve around two principal frameworks: multiple-strategy models and single-process models. Multiple-strategy models, often described using the "adaptive toolbox" metaphor, propose that individuals possess a repertoire of adaptive strategies. For instance, the Take-the-best heuristic distinguishes options based on the most important criterion. Another example of heuristic, not used to describe participant strategy in this task is the recognition heuristic. It involves considering the option we recognise as the more correct answer among the options (Gigerenzer & Todd, 1999). Single-process models, represented by the "adjustable spanner" metaphor (Newell, 2005), posit the existence of a uniform mechanism that is modulated based on the context (Glöckner & Betsch, 2008). These frameworks coexist due to their ability to explain empirical data but remain challenging to distinguish empirically (Newell & Bröder, 2008). More recent approaches, such as information intrusion paradigms, aim to differentiate these models by examining behaviours like visual search duration or stopping rules for visual information acquisition (Söllner et al., 2014). These studies reveal complex

interactions between cognitive costs, information structure, and contextual demands, highlighting that decisions result from a trade-off between effort and accuracy (Rieskamp & Otto, 2006).

### **Comparison with the set-selection paradigm**

To our knowledge, set selection has not been studied within the MADM paradigm. However, the paradigm could readily be adapted to set selection by increasing the number of options to select in each trial. Studies have shown that individuals can rapidly adjust their visual information exploration strategies as the number of available alternatives increases (Meißner et al., 2020). Specifically, they enhance their ability to filter information, avoiding exploration of less relevant details (Pieters & Warlop, 1999). Observing whether individuals can similarly adapt to the number of options to be selected in MADM tasks would be both pertinent and insightful.

The primary difference compared to our paradigm lies in information accessibility. In MADM, information is often abundant and readily available, meaning that information search resembles an integration process of varying depth. The paradigm assumes a world where all information is accessible but not necessarily integrable, whereas our paradigm operates in a world where accessing all information is impossible and must instead be guided by the selection objective.

## **V. Various sources of bias and mistakes**

We now present key concepts identified in the literature that may influence information search, decision-making under uncertainty, and the consideration of multiple choices. Unlike the previous section, the concepts presented here are not linked to a specific paradigm. Each bias or cognitive limitation is defined, illustrated with examples, and its connection to our paradigm is explained.

### **a. Confirmation biases**

Among cognitive biases, confirmation bias is particularly relevant to information processing (Lord et al., 1979; Nickerson, 1998; Wason, 1960, 1968). The literature identifies three forms of confirmation bias: (1) The tendency to seek information that supports one's hypothesis rather than information that might refute it, (2) The tendency to better recall information that confirms one's hypothesis compared to information that contradicts it, and (3) The tendency to interpret confirming information more accurately than disconfirming

information (Vedejová & Čavojová, 2021; Nickerson, 1998; Klayman, 1995). Note that the first form of confirmation bias has been described as a heuristic strategy of positive testing (Klayman & Ha, 1987) or matching bias (Evans, 1998).

One classic experimental paradigm used to study this phenomenon is the "2-4-6 task" by Wason (1960), where participants attempt to deduce a numerical rule by generating number triples. Most participants propose confirming triples, such as "8-10-12" for a hypothesis of adding two, rather than disconfirming triples, such as "5-5-6," which would provide more diagnostic information about the rule. This task illustrates the human tendency to rely on confirmatory evidence.

A similar confirmation bias appears in professional domains. Snyder & Swann (1978) trait hypothesis testing task demonstrated that individuals acting as interviewers preferred questions confirming their hypothesis about an interviewee's personality (e.g., testing extroversion with "What kinds of situations do you seek to meet new people?") rather than neutral or falsifying questions. This positive testing approach extends to expert contexts, such as judicial decision-making, where judges, despite their training, may exhibit biased information search when evaluating evidence. Research shows that judges can unintentionally overvalue evidence confirming their initial assessment of a case, leading to errors in judgement (Lallement et al., 2020).

Understanding the mechanisms and manifestations of confirmation bias, such as in heuristic-driven evidence search, highlights the need for interventions that encourage balanced and diagnostic information evaluation (Dickinson & Kakoschke, 2021).

In Chapter 2 of the thesis, we explore the potential links between the strategies employed by individuals in our paradigm and various tasks assessing confirmation bias. We seek to determine whether individuals' tendency to prioritise exploring their preferred options during the selection phase stems from confirmation bias.

### b. Pavlovian Biases

Biases in information sampling and decision making may be partly due to a Pavlovian behaviour. If an action is repeatedly associated with a reward, we tend to perform that action more easily than another (Hunt et al., 2016).

Hunt et al. study examines how participants make information-sampling decisions based on potential rewards. Participants play a game in which they pay to acquire information before selecting one of two options. Here, options are rows of cards (each composed of two cards) and depending on the trial, the row with the largest sum, the smallest sum, the smallest product, or the largest product is defined as the winning option. At the beginning of each trial, one of the four cards is revealed. Participants can pay 10 points to reveal a card (chosen by the game) or attempt to guess which row is the winner based solely on the available information. A correct guess earns 60 points, but an incorrect guess results in a loss of 50 points. If participants choose to reveal the proposed card, they may again decide between guessing the winning row immediately or paying an additional 15 points to reveal another card. This process can be repeated, with a final option to pay 20 points to reveal the last card.

The study identifies three main biases in information-sampling strategies. First, the "Sampling the Favourite" bias refers to participants' preference for sampling the row they believe is most likely to win, rather than the other row, which could provide more informative insights, as revealed in analyses across trials. Second, the "Positive Evidence Approach" reflects a tendency to stop sampling after revealing a card that confirms the preferred option, even when continuing to sample would reduce uncertainty. Third, the "Rejecting Unsampled Options" bias indicates that when participants refuse to explore a proposed card, they subsequently tend to avoid choosing the corresponding row. The authors interpret these biases as adaptive behaviours, possibly rooted in evolutionary perspectives. The tendency to sample based on preferred options may have been advantageous in natural decision-making contexts, where approaching potential sources of reward could maximise benefits.

We may question the connection between the bias of exploring preferred options described in this thesis and the Pavlovian bias "Sampling the Favourite" described by Hunt. Similarly, the bias "Rejecting the Unsampled Option" appears to describe behaviour closely related to the selection bias towards explored options that we outline. However, the significant differences between our two paradigms call for caution when interpreting the links between these biases (see further details in the General Discussion).

### c. Uncertainty Aversion

Uncertainty aversion originates from the work of Ellsberg (1961), who demonstrated that individuals prefer options with known probabilities over those with ambiguous probabilities, even when the objective outcomes are similar. This bias, linked to a negative perception of

uncertainty, manifests when individuals actively avoid choices where crucial information is missing or ambiguous.

A key example illustrating this bias is the "Ellsberg Paradox." Participants choose between two urns from which to draw a ball: one contains exactly 50 red balls and 50 black balls (known probabilities), while the other contains 100 balls whose colours (red or black) are in unknown proportions (ambiguous probabilities). Despite identical expected probabilities, participants consistently prefer the urn with known probabilities, highlighting their aversion to ambiguity. Studies in behavioural finance and management show that this bias reduces innovation and adoption of new technologies when success probabilities are not explicitly defined (Curley & Yates, 1989).

Uncertainty aversion does not appear to be directly observable in our paradigm, as the uncertainty surrounding unexplored options also leads to uncertainty about whether the explored options are the correct choices. We will revisit this in the General Discussion, where we will explore experimental possibilities for clarifying the effect of uncertainty on unexplored scores in participants' selection strategies.

### d. Choice Bracketing

The concept of choice bracketing describes how individuals group decisions into sets to evaluate their consequences. It distinguishes between narrow bracketing—where choices are considered in isolation—and broad bracketing—where multiple decisions are evaluated simultaneously to assess all effects. Choice bracketing is often influenced by the temporal sequencing of decisions. In narrow bracketing, choices are typically made one after the other. When encountering the first decision, an individual has no knowledge of the second decision they will need to make. In such situations, broad bracketing enables both choices to be considered simultaneously (Read et al., 2000).

An illustrative example of choice bracketing comes from an experiment in which participants faced two independent monetary decisions involving gains and losses. Two simultaneous financial decisions are presented in the task. In the first, participants choose between: (A) a certain gain of \$240 or (B) a 25% chance to win \$1,000 and a 75% chance to win nothing. In the second, they choose between: (C) a certain loss of \$750 or (D) a 75% chance to lose \$1,000 and a 25% chance to lose nothing. Most participants select option A (certain gain) in the first decision and option D (risk-taking) in the second. However, when explicitly asked to choose among combined options, no one selected the combination A + D, which is dominated by B +

C. It highlights that individuals tend to treat decisions in isolation rather than adopting a global perspective (Tversky & Kahneman, 1981). When presented with these decisions in isolation (narrow bracketing), participants often made globally inconsistent choices, such as accepting a certain gain in one scenario while taking a risk in the other. However, when the decisions were combined (broad bracketing), participants favoured a globally optimal option, highlighting the cumulative impacts of decisions (Read et al., 2000). Consequently, investors tend to exhibit myopic behaviour when frequently evaluating their portfolios, favouring low-risk, low-yield assets. In contrast, a broad bracketing framework encourages more rational choices by considering long-term returns (Benartzi & Thaler, 1995). Choice Bracketing may be related to the following concept of Contingent Thinking.

### e. Failure of Contingent Thinking (FCT)

A concept closely related to choice bracketing is contingent thinking. The "failure of contingent thinking" (FCT) revolves around individuals' inability to reason correctly about the hypothetical consequences of their actions in uncertain situations (Esponda & Vespa, 2014; Niederle & Vespa, 2023).

An illustrative example of FCT is observed in committee voting experiments, a classic decision-making scenario. In the experiment, an urn contains 10 balls: 3 blue and 7 red. A ball is drawn randomly, and participants act as members of a three-person committee tasked with determining the ball's colour through a majority vote. Two other "committee members" are computers following simple rules: one always votes for "red", and the other always votes correctly. The human participant must cast their vote, knowing the outcome depends on the majority. The optimal strategy is to determine when their vote is pivotal. If the ball is red, both computers vote red, guaranteeing a red majority. However, if the ball is blue, the computers' votes differ, making the participant's vote pivotal. Therefore, even with  $p(\text{blue}) = .3$ , the optimal strategy is always to vote blue. Yet, results show that only 15% of participants behave optimally (Esponda & Vespa, 2023).

In the basic condition, participants vote simultaneously with the computers, making it challenging to perceive the pivotal role of their vote. Esponda & Vespa (2014) implemented a remediation by conducting the vote sequentially rather than simultaneously. Performance improved significantly: around 75% of participants voted correctly when pivotal and voted indifferently when not. A further remediation involved informing participants that their vote would only be used if the computers disagreed, increasing the proportion of optimal

participants to 54%. These findings demonstrate that explicitly highlighting relevant contingencies reduces the effects of FCT to varying degrees.

In Chapter 4 of this thesis, we design an educational intervention aimed at remediation to help participants adopt a strategy that accounts for the relevant contingencies needed to solve the task. This involves explaining how to anticipate the use of information acquired during the exploration phase when making decisions in the selection phase.

## VI. Structure of the Thesis

This final section provides an overview of the chapters that structure the thesis, summarising their contributions to the understanding of information search strategies in set selection. Each chapter is dedicated to addressing specific aspects of the research questions introduced earlier.

### a. Chapter 1: Exploration Bias Towards Favourites in Set Selection

This chapter investigates how individuals approach information search when selecting a set of options. We compare participants' exploration strategy to the optimal and to several simple strategies. The study identifies two major biases: an exploration bias, where participants disproportionately examine options perceived as favourites for the final selection, and a selection bias, where explored options are favoured even when suboptimal. Through three experiments, the chapter demonstrates how these biases undermine performance.

### b. Chapter 2: Cognitive Factors Influencing Information Search Biases

Chapter 2 investigate the cognitive underpinnings of the exploration and selection biases observed in the first chapter. To understand the inter-individual variability of the biases observed in our paradigm, we use a set of tasks to serve as points of comparison.

Chapter 2 evaluates the relationships between these biases and traditional measures of cognitive abilities (Cognitive Reflection Test, Berlin Numeracy Test) and biases (e.g., confirmation biases, anchoring bias, framing effect, conjunction fallacy, base rate neglect, risk aversion). Despite scores on bias and ability tasks that are consistent with the literature, the findings reveal limited associations with exploration and selection biases in our paradigm.

As our task is inspired by sourcing and presented as a recruitment experiment, we wonder whether experts in the task present the same biases as non-experts in this type of task. We recruit a panel of professional recruiters via LinkedIn to run a shortened version of the experiment. We compare their answers with those of a non-expert population. We find that expertise does not necessarily result in the adoption of more optimal exploration strategies.

### c. Chapter 3: Metacognitive Judgements in Exploration Strategies for Set Selection

Based on the results of Chapters 1 and 2, individuals' strategy is stable across trials, and individuals can identify the heuristic that best describes their strategy. We therefore wonder whether they can evaluate the effectiveness of their strategy.

Chapter 3 examines how individuals perceive and evaluate their own information search strategies. It reveals a tendency for participants to overestimate the effectiveness of their preferred approaches, even when these are demonstrably suboptimal. However, confidence increases when participants successfully apply their chosen strategy. The findings highlight the role of metacognition in shaping exploration behaviours and decision confidence.

### d. Chapter 4: Learning and Adapting Optimal Strategies

As Chapter 3 shows that individuals overestimate their effectiveness, the question arises as to how best to help them. Among the possible means of remediation, educational remediation seems to be the best option when it can be implemented. Financial incentives can also help motivate participants to stay focused and avoid miscalculations.

Chapter 4 explores the capacity for learning and adaptation in the two stages of the paradigm. It shows that individuals can improve their exploration and selection strategies when provided with educative content supported by a training with feedback. Educated participants tends to generalize their knowledge to different versions of the tasks. The chapter also investigates the impact of financial incentives, finding no significant effect on strategy or performance.

## Bibliography

1. Abbas, S. I., Shah, M. H., & Othman, Y. H. (2021). Critical Review of Recruitment and Selection Methods: Understanding the Current Practices. *Annals of Contemporary Developments in Management & HR (ACDMHR)*, 3(3), Article 3. <https://doi.org/10.33166/ACDMHR.2021.03.005>
2. About Us. (n.d.). Retrieved 12 September 2024, from <https://news.linkedin.com/about-us>
3. Addicott, M. A., Pearson, J. M., Sweitzer, M. M., Barack, D. L., & Platt, M. L. (2017). A Primer on Foraging and the Explore/Exploit Trade-Off for Psychiatry Research. *Neuropsychopharmacology*, 42(10), 1931–1939. <https://doi.org/10.1038/npp.2017.108>
4. Bartoš, V., Bauer, M., Chytilová, J., & Matějka, F. (2016). Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition. *American Economic Review*, 106(6), 1437–1475. <https://doi.org/10.1257/aer.2014.0571>
5. Bellman, R. (1956). A Problem in the Sequential Design of Experiments. Indian Statistical Institute. <https://www.jstor.org/stable/25048278>
6. Bellman, R., & Kalaba, R. (1957). DYNAMIC PROGRAMMING AND STATISTICAL COMMUNICATION THEORY. *Proceedings of the National Academy of Sciences*, 43(8), 749–751. <https://doi.org/10.1073/pnas.43.8.749>
7. Benartzi, S., & Thaler, R. H. (1995). Myopic Loss Aversion and the Equity Premium Puzzle\*. *The Quarterly Journal of Economics*, 110(1), 73–92. <https://doi.org/10.2307/2118511>
8. Bertrand, M., & Duflo, E. (2017). Field Experiments on Discrimination. In A. V. Banerjee & E. Duflo (Eds.), *Handbook of Economic Field Experiments* (Vol. 1, pp. 309–393). North-Holland. <https://doi.org/10.1016/bs.hefe.2016.08.004>
9. Bordalo, P., Coffman, K., Gennaioli, N., & Shleifer, A. (2016). Stereotypes\*. *The Quarterly Journal of Economics*, 131(4), 1753–1794. <https://doi.org/10.1093/qje/qjw029>
10. Cauthen, N. R., Robinson, I. E., & Krauss, H. H. (1971). Stereotypes: A Review of the Literature 1926–1968. *The Journal of Social Psychology*. <https://www.tandfonline.com/doi/abs/10.1080/00224545.1971.9918526>
11. Chow, Y. S., Moriguti, S., Robbins, H., & Samuels, S. M. (1964). Optimal selection based on relative rank (the “secretary problem”). *Israel Journal of Mathematics*, 2(2), 81–90. <https://doi.org/10.1007/BF02759948>
12. Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933–942. <https://doi.org/10.1098/rstb.2007.2098>
13. Cole, M. S., Rubin, R. S., Feild, H. S., & Giles, W. F. (2007). Recruiters’ Perceptions and Use of Applicant Résumé Information: Screening the Recent Graduate. *Applied Psychology*, 56(2), 319–343. <https://doi.org/10.1111/j.1464-0597.2007.00288.x>
14. Corbin, R. M. (1980). The secretary problem as a model of choice. *Journal of Mathematical Psychology*, 21(1), 1–29. [https://doi.org/10.1016/0022-2496\(80\)90025-5](https://doi.org/10.1016/0022-2496(80)90025-5)
15. Curley, S. P., & Yates, J. F. (1989). An empirical evaluation of descriptive models of ambiguity reactions in choice situations. *Journal of Mathematical Psychology*, 33(4), 397–427. [https://doi.org/10.1016/0022-2496\(89\)90019-9](https://doi.org/10.1016/0022-2496(89)90019-9)
16. Daw, N. D., O’Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876–879. <https://doi.org/10.1038/nature04766>
17. Dickinson, D. L., & Kakoschke, N. (2021). Seeking confirmation? Biased information search and deliberation in the food domain. *Food Quality and Preference*, 91, 104189. <https://doi.org/10.1016/j.foodqual.2021.104189>
18. Drezner, W., & Edigbe, E. (2024). Accessible Low-Code No-Code Development. <https://www.diva-portal.org/smash/get/diva2:1879862/FULLTEXT01.pdf>
19. D’Silva, C. (2020). A Study On Increase in E-Recruitment and Selection Process. *International Journal of Research in Engineering, Science and Management*, 3(8), Article 8.

20. Ellemers, N. (2018). Gender Stereotypes. *Annual Review of Psychology*, 69(Volume 69, 2018), 275–298. <https://doi.org/10.1146/annurev-psych-122216-011719>
21. Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms\*. *The Quarterly Journal of Economics*, 75(4), 643–669. <https://doi.org/10.2307/1884324>
22. Esponda, I., & Vespa, E. (2014). Hypothetical Thinking and Information Extraction in the Laboratory. *American Economic Journal: Microeconomics*, 6(4), 180–202. <https://doi.org/10.1257/mic.6.4.180>
23. Esponda, I., & Vespa, E. (2023). Contingent Thinking and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory. *The Review of Economic Studies*, rdad102. <https://doi.org/10.1093/restud/rdad102>
24. Evans, J. St. B. T. (1998). Matching Bias in Conditional Reasoning: Do We Understand it After 25 Years? *Thinking & Reasoning*, 4(1), 45–110. <https://doi.org/10.1080/135467898394247>
25. Eyting, M. (2022). Why do we Discriminate? The Role of Motivated Reasoning. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4210315>
26. Ferguson, T. S. (1989). Who Solved the Secretary Problem? *Statistical Science*, 4(3), 282–289. <https://doi.org/10.1214/ss/1177012493>
27. Freeman, P. R. (1983). The Secretary Problem and Its Extensions: A Review. *International Statistical Review / Revue Internationale de Statistique*, 51(2), 189–206. <https://doi.org/10.2307/1402748>
28. Gianini-Pettit, J. (1979). Optimal Selection Based on Relative Ranks with a Random Number of Individuals. *Advances in Applied Probability*, 11(4), 720–736. <https://doi.org/10.2307/1426856>
29. Gigerenzer, G. (2024). The rationality wars: A personal reflection. *Behavioural Public Policy*, 1–21. <https://doi.org/10.1017/bpp.2024.51>
30. Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98(4), 506–528. <https://doi.org/10.1037/0033-295X.98.4.506>
31. Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart* (pp. 3–34). Oxford University Press.
32. Gittins, J. C. (1975). The Two-Armed Bandit Problem: Variations on a Conjecture by H. Chernoff.
33. Gittins, J. C. (1979). Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 41(2), 148–177.
34. Gittins, J. C., & Jones, D. M. (1979). A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66(3), 561–565. <https://doi.org/10.1093/biomet/66.3.561>
35. Glöckner, A., & Betsch, T. (2008). Multiple-reason decision making based on automatic processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(5), 1055–1075. <https://doi.org/10.1037/0278-7393.34.5.1055>
36. Hausmann, D., & Läge, D. (2008). Sequential evidence accumulation in decision making: The individual desired level of confidence can explain the extent of information acquisition. *Judgment and Decision Making*, 3(3), 229–243. <https://doi.org/10.1017/S1930297500002436>
37. Hoffman, K. L., & Padberg, M. (2001). Traveling salesman problem. In S. I. Gass & C. M. Harris (Eds.), *Encyclopedia of Operations Research and Management Science* (pp. 849–853). Springer US. [https://doi.org/10.1007/1-4020-0611-X\\_1068](https://doi.org/10.1007/1-4020-0611-X_1068)
38. Hosain, M. S., & Liu, P. (2020). The Role of Social Media on Talent Search and Acquisition: Evidence from Contemporary Literature. *Journal of Intercultural Management*, 12(1), 92–137. <https://doi.org/10.2478/joim-2020-0034>
39. Hunt, L. T., Rutledge, R. B., Malalasekera, W. M. N., Kennerley, S. W., & Dolan, R. J. (2016). Approach-Induced Biases in Human Information Sampling. *PLOS Biology*, 14(11), e2000638. <https://doi.org/10.1371/journal.pbio.2000638>
40. Jowell, R., & Prescott-Clarke, P. (1970). Racial Discrimination and White-collar Workers in Britain. *Race*, 11(4), 397–417. <https://doi.org/10.1177/030639687001100401>

41. Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research*, 4, 237–285. <https://doi.org/10.1613/jair.301>
42. Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237.
43. Kelly, Christopher. A., & Sharot, T. (2021). Individual differences in information-seeking. *Nature Communications*, 12(1), 7062. <https://doi.org/10.1038/s41467-021-27046-5>
44. Klayman, J. (1995). Varieties of Confirmation Bias. In J. Busemeyer, R. Hastie, & D. L. Medin (Eds.), *Psychology of Learning and Motivation* (Vol. 32, pp. 385–418). Academic Press. [https://doi.org/10.1016/S0079-7421\(08\)60315-1](https://doi.org/10.1016/S0079-7421(08)60315-1)
45. Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94(2), 211–228. <https://doi.org/10.1037/0033-295X.94.2.211>
46. Kleinberg, R. (2005). A multiple-choice secretary algorithm with applications to online auctions. *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 630–631.
47. Kobayashi, K., & Hsu, M. (2019). Common neural code for reward and information value. *Proceedings of the National Academy of Sciences*, 116(26), 13061–13066. <https://doi.org/10.1073/pnas.1820145116>
48. Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature Human Behaviour*, 3(6), 587–595. <https://doi.org/10.1038/s41562-019-0589-3>
49. Koch, A. J., D’Mello, S. D., & Sackett, P. R. (2015). A meta-analysis of gender stereotypes and bias in experimental simulations of employment decision making. *Journal of Applied Psychology*, 100(1), 128–161. <https://doi.org/10.1037/a0036734>
50. Kroll, E., Veit, S., & Ziegler, M. (2021). The Discriminatory Potential of Modern Recruitment Trends—A Mixed-Method Study From Germany. *Frontiers in Psychology*, 12. <https://www.frontiersin.org/articles/10.3389/fpsyg.2021.634376>
51. Lallement, J., Dejean, S., Euzéby, F., & Martinez, C. (2020). The interaction between reputation and information search: Evidence of information avoidance and confirmation bias. *Journal of Retailing and Consumer Services*, 53, 101787. <https://doi.org/10.1016/j.jretconser.2019.03.014>
52. Lang, K., & Lehmann, J.-Y. K. (2012). Racial Discrimination in the Labor Market: Theory and Empirics. *Journal of Economic Literature*, 50(4), 959–1006. <https://doi.org/10.1257/jel.50.4.959>
53. Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the “take the best” and the “rational” models. *Psychonomic Bulletin & Review*, 11(2), 343–352. <https://doi.org/10.3758/BF03196581>
54. Lindley, D. V. (1961). Dynamic Programming and Decision Theory. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 10(1), 39–51. <https://doi.org/10.2307/2985407>
55. LinkedIn recruiting trends (fr). (n.d.). Retrieved 14 September 2024, from [https://business.linkedin.com/content/dam/business/talent-solutions/regional/fr\\_FR/site/pdf/playbooks/linkedin-recruiting-trends-fr.pdf](https://business.linkedin.com/content/dam/business/talent-solutions/regional/fr_FR/site/pdf/playbooks/linkedin-recruiting-trends-fr.pdf)
56. Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11), 2098–2109. <https://doi.org/10.1037/0022-3514.37.11.2098>
57. Meißner, M., Oppewal, H., & Huber, J. (2020). Surprising adaptivity to set size changes in multi-attribute repeated choice tasks. *Journal of Business Research*, 111, 163–175. <https://doi.org/10.1016/j.jbusres.2019.01.008>
58. Mucci, A. G. (1973). On a Class of Secretary Problems. *The Annals of Probability*, 1(3), 417–427.
59. Neumark, D. (2018). Experimental Research on Labor Market Discrimination. *Journal of Economic Literature*, 56(3), 799–866. <https://doi.org/10.1257/jel.20161309>
60. Newell, B. R. (2005). Re-visions of rationality? *Trends in Cognitive Sciences*, 9(1), 11–15. <https://doi.org/10.1016/j.tics.2004.11.005>
61. Newell, B. R., & Bröder, A. (2008). Cognitive processes, models and metaphors in decision research. *Judgment and Decision Making*, 3(3), 195–204. <https://doi.org/10.1017/S1930297500002400>

62. Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of "success". *Journal of Behavioral Decision Making*, 17(2), 117–137. <https://doi.org/10.1002/bdm.465>
63. Nickerson, R. S. (1998). Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology*, 2(2), 175–220. <https://doi.org/10.1037/1089-2680.2.2.175>
64. Niederle, M., & Vespa, E. (2023). Cognitive Limitations: Failures of Contingent Thinking. *Annual Review of Economics*, 15(Volume 15, 2023), 307–328. <https://doi.org/10.1146/annurev-economics-091622-124733>
65. Pager, D., Bonikowski, B., & Western, B. (2009). Discrimination in a Low-Wage Labor Market: A Field Experiment. *American Sociological Review*, 74(5), 777–799. <https://doi.org/10.1177/000312240907400505>
66. Petrucci, J. D. (1981). Best-Choice Problems Involving Uncertainty of Selection and Recall of Observations. *Journal of Applied Probability*, 18(2), 415–425. <https://doi.org/10.2307/3213287>
67. Pieters, R., & Warlop, L. (1999). Visual attention during brand choice: The impact of time pressure and task motivation. *International Journal of Research in Marketing*, 16(1), 1–16. [https://doi.org/10.1016/S0167-8116\(98\)00022-6](https://doi.org/10.1016/S0167-8116(98)00022-6)
68. Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, 5(2), 105–120. [https://doi.org/10.1016/0030-5073\(70\)90008-5](https://doi.org/10.1016/0030-5073(70)90008-5)
69. Rasmussen, W. T., & Robbins, H. (1975). The Candidate Problem with Unknown Population Size. *Journal of Applied Probability*, 12(4), 692–701. <https://doi.org/10.2307/3212720>
70. Read, D., Loewenstein, G., Rabin, M., Keren, G., & Laibson, D. (2000). Choice Bracketing. In B. Fischhoff & C. F. Manski (Eds.), *Elicitation of Preferences* (pp. 171–202). Springer Netherlands. [https://doi.org/10.1007/978-94-017-1406-8\\_7](https://doi.org/10.1007/978-94-017-1406-8_7)
71. Rieskamp, J., & Otto, P. E. (2006). SSL: A Theory of How People Learn to Select Strategies. *Journal of Experimental Psychology: General*, 135(2), 207–236. <https://doi.org/10.1037/0096-3445.135.2.207>
72. Roulin, N., & Levashina, J. (2019). LinkedIn as a new selection method: Psychometric properties and assessment approach. *Personnel Psychology*, 72(2), 187–211. <https://doi.org/10.1111/peps.12296>
73. Savage, L. J. (1954). *The foundations of statistics* (pp. xv, 294). John Wiley & Sons.
74. Savage, L. J. (1972). *The Foundations of Statistics*. Courier Corporation.
75. Seale, D. A., & Rapoport, A. (1997). Sequential Decision Making with Relative Ranks: An Experimental Investigation of the 'Secretary Problem'. *Organizational Behavior and Human Decision Processes*, 69(3), 221–236. <https://doi.org/10.1006/obhd.1997.2683>
76. Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. *Nature Human Behaviour*, 4(1), 14–19. <https://doi.org/10.1038/s41562-019-0793-1>
77. Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99. <https://doi.org/10.2307/1884852>
78. Smith, M. H. (1975). A Secretary Problem with Uncertain Employment. *Journal of Applied Probability*, 12(3), 620–624. <https://doi.org/10.2307/3212880>
79. Snyder, M., & Swann, W. B. (1978). Hypothesis-testing processes in social interaction. *Journal of Personality and Social Psychology*, 36(11), 1202–1212. <https://doi.org/10.1037/0022-3514.36.11.1202>
80. Söllner, A., Bröder, A., Glöckner, A., & Betsch, T. (2014). Single-process versus multiple-strategy models of decision making: Evidence from an information intrusion paradigm. *Acta Psychologica*, 146, 84–96. <https://doi.org/10.1016/j.actpsy.2013.12.007>
81. Söllner, A., Bröder, A., & Hilbig, B. E. (2013). Deliberation versus automaticity in decision making: Which presentation format features facilitate automatic decision making? *Judgment and Decision Making*, 8(3), 278–298. <https://doi.org/10.1017/S1930297500005982>
82. Statista. (2024). Number of e-mail users worldwide 2027. Statista. <https://www.statista.com/statistics/255080/number-of-e-mail-users-worldwide/>

## General introduction

83. Stewart, T. J. (1981). The Secretary Problem with an Unknown Number of Options. *Operations Research*, 29(1), 130–145. <https://doi.org/10.1287/opre.29.1.130>
84. Sue, S., & Kitano, H. H. L. (1973). Stereotypes as a Measure of Success. *Journal of Social Issues*, 29(2), 83–98. <https://doi.org/10.1111/j.1540-4560.1973.tb00074.x>
85. Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press.
86. Thrun, S. B. (1992). THE ROLE OF EXPLORATION IN LEARNING CONTROL.
87. Tversky, A., & Kahneman, D. (1981). The Framing of Decisions and the Psychology of Choice. *Science*, 211(4481), 453–458. <https://doi.org/10.1126/science.7455683>
88. Wason, P. C. (1960). On the Failure to Eliminate Hypotheses in a Conceptual Task. *Quarterly Journal of Experimental Psychology*, 12(3), 129–140. <https://doi.org/10.1080/17470216008416717>
89. Wason, P. C. (1968). Reasoning about a Rule. *Quarterly Journal of Experimental Psychology*, 20(3), 273–281. <https://doi.org/10.1080/14640746808400161>
90. Weber, R. (1992). On the Gittins Index for Multiarmed Bandits. *The Annals of Applied Probability*, 2(4). <https://doi.org/10.1214/aoap/1177005588>

# Chapter 1: Exploration Bias Towards Favourites in Set Selection

Jean-Michel Dagba<sup>1,2</sup>, Jean-Christophe Vergnaud<sup>1,3</sup>, Vincent de Gardelle<sup>1,3,4</sup>

<sup>1</sup> Centre d'Économie de la Sorbonne & Université Paris 1 Panthéon Sorbonne, France

<sup>2</sup> Entreprise Humans Matter, France

<sup>3</sup> Centre national de la recherche scientifique, France

<sup>4</sup> Paris School of Economics, France

## Abstract

The present study investigates information search strategies when individuals must select not the single best option in an ensemble but a set of best options, much like a student choosing which universities to apply to, or a university selecting the best students. In such tasks, exploring the a priori favourite options seems intuitive, but this strategy is in fact suboptimal. We report three experiments employing an original two-stage task, where participants (total N = 196) first reveal the hidden scores for some options (exploration stage) and then pick the best options (selection stage). Our results indicate suboptimal performance in all experiments, characterized by a bias towards the favourites options during the exploration stage (with some inter-individual variability), and a bias towards the explored options during the selection stage. We discuss the implications of these results for real life situations, such as recruitment procedures.

## Research Transparency Statement

Conflicts of interest: All authors declare no conflicts of interest. Funding: This research was supported by Centre d'Économie de la Sorbonne of Université Panthéon-Sorbonne Paris 1 and the company Humans Matter. Artificial intelligence: ChatGPT was use for the writing of the manuscript to improve the translation from French to English. Ethics: This research received approval from the Institutional Review Board of Paris School of Economics (decision 2021-024). Computational reproducibility: All study materials, primary data and analysis scripts for all experiments reported in this manuscript are publicly available ([project link on OSF](#)). Preregistration: No aspects of the study were preregistered.

## Table of contents

Abstract	32
Research Transparency Statement	32
Table of contents	33
I. Introduction	35
II. General method	36
a. Participants	36
b. Apparatus	36
c. Stimuli and procedure	36
d. Model and measures	38
e. Statistical analyses	40
III. Experiment 1	41
a. Method	41
b. Results	41
IV. Experiment 2	45
a. Motivation	45
b. Method	46
c. Results	47
V. Experiment 3	50
a. Motivation	50
b. Method	50
c. Results	50
VI. Discussion	53
Bibliography	56
Supplementary materials	59
Expected performance, optimal selection and optimal exploration	59



## I. Introduction

In the decision process, searching for information before committing to a choice is central. For example, a student who wishes to determine where to apply may first consider the course contents and prestige of several universities, and then she would visit the campus of her favourite option to confirm it. But what if the student must indicate a wish-list of 5 universities? Clearly, visiting all the campuses will be too costly, so the student may prefer not visit universities she is already happy about, given the other information she has. Alternatively, she may still like to confirm the quality of her favourite universities. So what would it be? So far, the literature on information search has mainly focused on the selection of a single option, leaving aside situations in which individuals must select a set of several options. Our goal in the current paper is to start investigating this issue, by designing an original experimental task and documenting how individuals may or may not find the optimal information search strategy in this task.

Prior research has shown that when they have to pick a single option, individuals are not optimal in their search for information. For instance, they might sample too much (Juni et al., 2016) or under-react to manipulations of information costs (Bouhlel et al., 2022). More generally, search behaviour has been described as based on heuristics (Gigerenzer & Gaissmaier, 2011; Lee & Cummins, 2004; Newell et al., 2004), in line with the influential idea that humans do not necessarily aim for optimal solutions, which can be too costly to compute, but may be satisfied with strategies offering a lower but “good-enough” level of performance (Simon, 1955) or confidence (Hausmann & Läge, 2008). This need for confidence may explain why individuals seek non-instrumental information (Eliaz & Schotter, 2007, 2010; Kobayashi et al., 2019; Matthews et al., 2023), or exhibit confirmation biases when searching for information (Klayman, 1995; Rassin et al., 2010). Indeed, such biases would allow individuals to accumulate rapidly evidence in favour of their preferred option, and to increase confidence in their choice.

In the context of set choices, such heuristics or biases are presumably even more likely to operate, because figuring out the optimal strategy is computationally more complex. However, although the heuristic may be a reasonable approximation of the optimal choice in single choice situation, it is not so clear that it is the case when choosing a set of options. In our introductory example, visiting the 5 campuses at the top of the list is the optimal strategy if the goal is to pick the best university, but not if our goal is to establish a wish list of 5 universities, as we shall see below.

To investigate how individuals search for information when having to select a set of options, we designed an original task in 2 stages. The task was framed as a recruitment situation and participants faced 10 options (candidates) defined by two scores, one known and one hidden (Fig. 1a). After revealing the hidden score for some options (exploration stage), participants must identify the best options overall (selection stage, Fig. 1b). In this setting, we can illustrate that having to pick not just one but several options at the selection stage has consequences for the optimal exploration strategy (Fig. 1c). To characterize participants' behaviour, we compare it to the optimal strategy and to several exploration heuristics followed by an optimal selection according to the revealed information (Fig. 1d). Experiment 1 is our baseline condition with 5 options in the exploration stage and 5 options in the selection stage. In Experiments 2 and 3 we vary the size of the exploration set, of the selection set, and the framing of the task respectively. In all experiments, we found that most participants focus too much on the best options during the exploration stage and exhibit a bias towards explored options at the selection stage, failing to adopt a simple heuristic that closely approximates optimal performance.

## II. General method

### a. Participants

196 healthy French adults (100 women), with a mean age of 32.34 years-old, took part in our experiments (around 50 participants per experiment). They were recruited from the Laboratoire d'Economie Expérimentale de Paris pool. They provided consent to take part in an experiment of 45 minutes.

### b. Apparatus

Experiments were conducted online. They were build using oTree (Chen et al., 2016) with the technical support of the Fédération S2CH.

### c. Stimuli and procedure

Participants' main task was framed as a recruitment task: on each trial, they were presented with 10 profiles, and they had to find the best 5 out of these 10 profiles. Each profile consisted of 2 scores, noted A and B. All scores were drawn independently from a uniform distribution between 0 and 10, and this was known to participants. At the beginning of the trial, all scores A are revealed but scores B are not. In the exploration stage, participants must pick 5 profiles for which the value of score B will be revealed. Then, in the selection stage, participants must

select the 5 best profiles, defined as the ones with the highest means of A and B. Participants were told that scores on A and B were equally important, even for scores that are not revealed. Following instructions, participants engaged in a training trial, in which all scores were revealed after the selection phase, to ensure that they understood the task and the importance of unrevealed information.

After each trial, participants received feedback about the performance of their choices, that is, about the number of profiles in their selection that belonged to the set of best profiles. Note that in case of equality, the set of best profiles can be larger than 5.

Each participant completed a total of 25 trials, in 25 to 30 minutes (including instructions). Participants could earn a bonus payment contingent on their performance: one trial was selected at random, and they received 1 euro for each correct answer in this trial.

Across our different experiments (see below), we manipulated some specifications of the main task: in Experiment 2A, the size of the exploration set was of 3 profiles instead of 5, in Experiment 2B the selection set was of 8 profiles, and in Experiment 4 we changed the framing of the selection stage from selecting the best to rejecting the worst profiles. Unless noted otherwise, all other parameters were identical.

After the main task, participants completed various test to measure confirmation bias, risk preferences and cognitive abilities (Berlin Numeracy Test and Cognitive Reflection Test), which we will not be analysed in the present work.

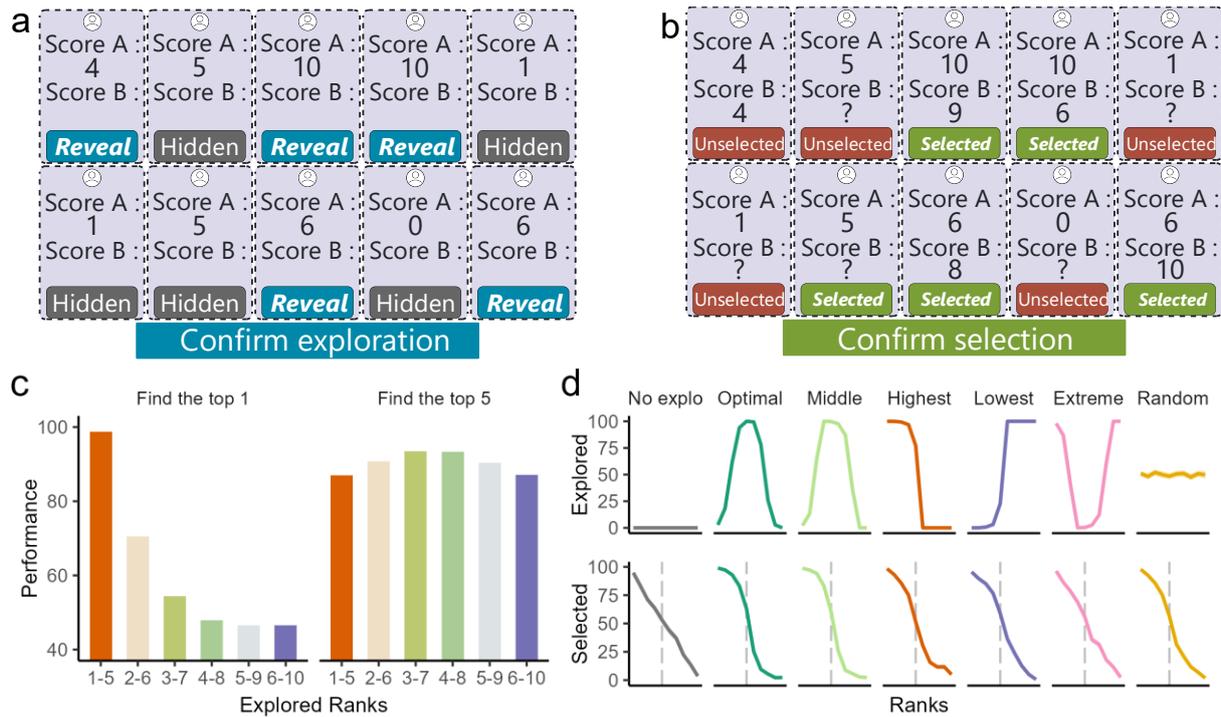


Fig 1.

Presentation of the main task and the consequences of set selection according to the strategy used. The two stages during the experiment are (a) exploration stage: participants consult the score A of the 10 options and indicates 5 options for which they wish to reveal scores B which are hidden at first, and (b) the selection stage: participants access the requested scores B and they must select the 5 best options from the group. How to explore when the goal is to select the single best option vs. the 5 best options? We calculated (c) the performance of a virtual agent exploring the 5 highest scores A, or scores with ranks 2-6, etc., and performing ideal selection based on this information. When the goal is to select the best option, we find that the optimal strategy is to explore ranks 1-5, whereas when the 5 best options must be found, it becomes optimal to explore intermediate ranks. Finally, we illustrated (d) various exploration strategies. The panels illustrate the proportion of profiles explored and selected, as a function of their ranks on score A (for the exploration phase, top panels) or on the total score (for the selection phase, bottom panels). No-explo: The agent receives no information during the exploration phase. Optimal: the agent performs optimal exploration on each trial. The other exploration strategies focus on the 5 top options on score A (Highest), the 5 bottom options (Lowest), the middle-ranked options (Middle), the best- and worst-ranked options (Extreme) or a random exploration of 5 options (Random). Each exploration strategy involves optimal selection given the explored information.

#### d. Model and measures

##### Calculation of performance

Performance for a given selection is defined as the fraction of selected options that verify the objective (e.g. top 5 in Exp. 1 and 2A, top 8 in Exp. 2B, bottom 5 in Exp. 3). Normalized performance was calculated as the performance of the participant minus the chance level, divided by the distance between optimal and chance-level performance.

## Strategies for exploration

To characterize participants' behaviour in the exploration stage, we compared their exploration responses to several heuristics: examining options with the highest ranks (*Highest*), with intermediate ranks (*Middle*), with the lowest ranks (*Lowest*) or with ranks at the high and low extremes (*Extreme*). Note that the *Middle* heuristic involves a random choice between the option ranked 3<sup>rd</sup> and the option ranked 8<sup>th</sup>, in addition to exploring all options ranked 4<sup>th</sup> to 7<sup>th</sup>. Similarly, for the *Extreme* heuristic, a choice is made between the option ranked 3<sup>rd</sup> and the one ranked 8<sup>th</sup>, in addition to explorations of options at the first and last two ranks. In addition, we considered a random exploration strategy (*Random*), a strategy with no exploration (*No Explo*), and we estimated the optimal exploration strategy (*Optimal*), as detailed below and in the Supplementary Material.

## Optimal selections conditional on exploration

When evaluating the performance of the different exploration strategies, we take into account how these exploration strategies may lead to different selections. Indeed, the best selection approach can depend on the information obtained during the exploration phase. We thus considered, for each exploration strategy, the selection that would maximize expected performance (for details see Supplementary Materials).

## Optimal exploration

This optimal exploration is defined as the exploration that leads to maximal expected performance, when it is followed by an optimal selection process. Despite the apparent simplicity of the paradigm, we could only estimate the optimal exploration strategy using a sampling approach (for details see Supplementary Materials). As illustrated in Fig. 1d, the optimal exploration strategy in the setting of Experiment 1 is close to the *Middle* heuristic. However, it is not exactly identical to the *Middle* heuristic, because it uses both the options' ranks and their numerical values. For instance, if there is a large gap in the score A between the 6<sup>th</sup> and the 8<sup>th</sup> options but a small gap between the 3<sup>rd</sup> and the 5<sup>th</sup> options, the optimal exploration will include the 3<sup>rd</sup> option but not the 8<sup>th</sup> one. In contrast, the *Middle* strategy considers only ranks, so it would randomly choose between the two during the exploration phase. In Fig. 1a the *Optimal* strategy is to explore options scores 4, 5, 5, 6 and 6, whereas the *Middle* strategy explores scores 4, 5, 5 and 6 (as the *Optimal* strategy does) but randomly picks between the second option scored 6 and the one scored 1.

The optimal strategy outlined here would consistently outperform any other heuristic in cases where scores A and B are correlated. Specifically, the Highest strategy is not optimal when the scores are correlated (see Figure 10 in the Supplementary Materials).

### Calculation of biases and errors

During the exploration phase, the bias towards the highest options was defined as the difference between the proportion of options with the highest scores A explored by participants, minus the same proportion evaluated for the optimal exploration strategy. A positive value indicates that participants explored the options with highest scores A more than prescribed by the optimal solution.

During the selection phase, a bias towards the explored options was defined as the difference between the proportion of explored options chosen by participants and the same proportion evaluated for the optimal solution (following participants exploration). A positive value indicates that participants exhibit a bias towards explored options during the selection phase. We also evaluated "mathematical errors" made by participants during their selection, separately for explored options and for unexplored options. To do so, we evaluated whether when selecting options within the explored (resp. unexplored) options, they actually picked the ones with the highest total scores (resp. the highest score A). We summed the errors made for explored and unexplored options and divided this sum by the total number of options selected, to express these errors as a percentage.

### e. Statistical analyses

For pairwise comparisons of performance (between participants, or between models) we used Wilcoxon rank test and we report the effect size as  $r$  and 95% confidence intervals. We used chi-square tests and report Adjusted Cramer's  $v$  as the effect size to compare the proportions of the best-matching strategy across experiments. For exploration times, at the individual level we used the median duration across trials and compared subgroups of participants using Wilcoxon rank tests. Statistical analyses were conducted using R (R Core Team, 2023), as well as the packages 'tidyverse' (Wickham et al., 2019), 'ggpubr' (Kassambara, 2023), 'cowplot' (O. Wilke, 2024), 'easystats' (Lüdtke et al., 2022), and 'vcd' (Zeileis et al., 2007).

### III. Experiment 1

#### a. Method

Experiment 1 serves as our baseline condition, featuring an exploration set of 5 options and a selection set of 5 options. In this setting, the different strategies produce distinct exploration and selection responses (see Fig. 1c and Supplementary Material), and the optimal strategy is close to the *Middle* heuristic.

A total of 49 participants (22 women) participated in Experiment 1, with a mean age of 33.6 years (ranging from 18 to 72, with a median age of 24 years). Their mean gain was 4.37€. This sample size was not based on prior studies, as the experimental procedure is novel.

#### b. Results

##### Performance

###### *Global performance*

We first examined participants' performances, i.e. whether they actually selected the 5 best profiles in each trial. As illustrated on Fig. 2a, most participants performed relatively well, with 83% (95% CI = [80.95%, 84.41%]) of correct selections on average, largely exceeding chance performance in this setting (chance-level 54%:  $W = 1125.00$ ,  $p < .001$ ;  $r = 1.00$ , 95% CI = [1.00, 1.00]). Participants' performance also exceeded that of a virtual agent with no exploration phase (78%, 95% CI = [76.69%, 78.83%];  $W = 949.00$ ,  $p < .001$ ;  $r = .83$ , 95% CI = [.71, .91]), indicating that they used the information they collected in the exploration stage. However, they failed short of optimal performance (92%, 95% CI = [91.07%, 92.61%],  $W = 1176.00$ ,  $p < .001$ ;  $r = 1.00$ , 95% CI = [1.00, 1.00]). Fig. 2b further illustrates this: in comparing to the optimal agent, participants selected too many options with low ranks and missed some of the highest options.

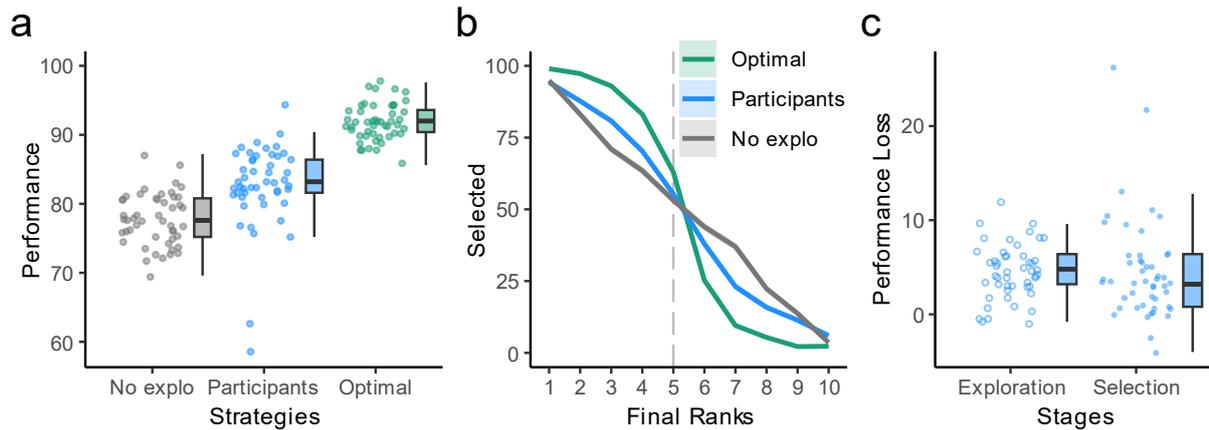


Fig 2.

Participants' performance and loss in comparison to the performance of the optimal agent (see methods for details) and to the performance of a virtual agent who receives no information during the exploration phase. These figures demonstrate (a) the performance in Experiment 1, as the proportions of selected options that were in the best set, and (b) proportions of selected options within each rank (averaged across participants), for participants, for the optimal, and the no-exploration agent. Confidence intervals for each curve are smaller than the curve width. Finally, (c) the performance loss (in percentage points) compared to the optimal agent was decomposed into exploration and selection components.

### Loss decomposition

We further decomposed this performance loss into two components: the loss due to suboptimal exploration and the loss due to suboptimal selection. To do so, we evaluated for each participant a virtual agent using an optimal selection strategy when given the information explored by the participant. This virtual agent achieved a performance of 87% on average (95% CI = [86.49, 88.14]), which indicates that because of poor selection, participants incurred a significant loss of 4.64 percentage points in their performance (95% CI = [3.09, 6.19],  $W = 955.00$ ,  $p < .001$ ;  $r = .93$ ). The performance of this virtual agent was 4.52 percentage points below that of the optimal agent, indicating a significant loss of performance also at the exploration stage (95% CI = [3.69, 5.36],  $W = 1162.00$ ,  $p < .001$ ;  $r = .98$ ). We further analyse each stage separately in the sections below.

## Exploration stage

### Optimal and heuristic strategies

Fig. 3a illustrates the profiles explored by participants, as a function of the rank of the score A. As we can see, participants focused on the highest scores, whereas the optimal exploration is centred on the intermediate ranks. We used the overlap between the items explored to quantify the similarity between participants' exploration strategy and various strategies (Fig. 3b). Participants' exploration was most similar to the *Highest* strategy (mean overlap: 76%, 95% CI = [68.17%, 83.73%]), in comparison to the *Optimal* (58%, 95% CI = [54.91%, 61.01%])

and *Middle* (57%, 95% CI = [54.03%, 60.19%]) strategies. This was confirmed when examining the best-matching strategy at the individual level: a large majority of participants (65%) were closest to the *Highest* heuristic (Fig. 3c). Following the idea that the *Highest* strategy was most intuitive for participants, we also examined how fast participants completed the exploration stage, and found that those who followed predominantly the *Highest* strategy were faster than those who did not (median completion time: 8952ms vs. 11594ms;  $W = 147.00$ ,  $p = .008$ ;  $r = -.46$ , 95% CI = [-.69, -.15]).

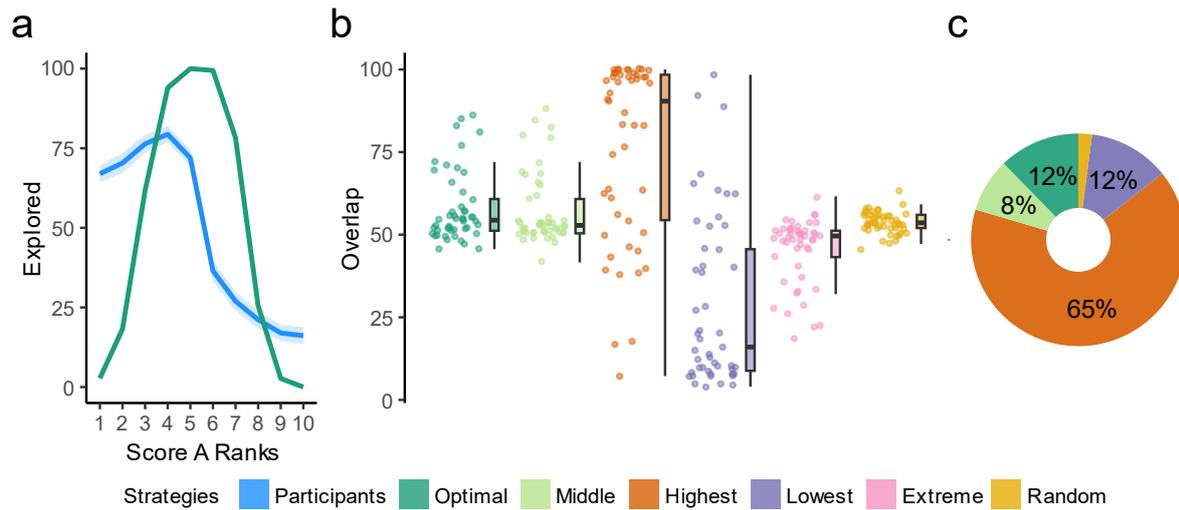


Fig. 3.

Participants' exploration compared with the optimal strategy and various heuristics in Experiment 1. Graphics represents (a) the proportions of explored profiles within each rank, for participants and for the optimal agent, (b) the proportions of overlap between participants exploration and optimal exploration / heuristics, and (c) proportions of the best-matching strategies, across participants.

### *Bias towards exploring the highest options*

As previously observed, participants over-explored the 5 options with the highest ranks on scores A (78%, 95% CI = [70.99%, 85.58%]) in comparison to the optimal agent who would explore these options much less (60%, 95% CI = [59.12%, 61.73%];  $W = 1013.50$ ,  $p < .001$ ;  $r = .65$ , 95% CI = [.43, .80]). In other words, we found evidence of a bias towards exploring options that are a priori the best options, given their score A.

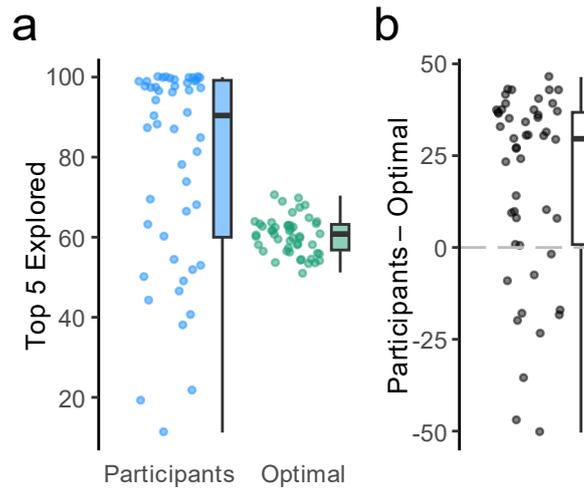


Fig 4.

Exploration bias toward the best options (i.e. options with a score A ranked between 1 and 5). We represented first (a) the proportions of best options explored by participants and by the optimal agent, and secondly (b) the difference of proportions between participant and optimal agent ‘best options’ exploration, showing that participants over-explored options with the highest scores A.

### Selection stage

To evaluate the quality of participants’ responses during the selection phase, we compared them with those of a virtual agent who would have obtained the same information during the exploration, and who would use this information optimally to maximise the expected performance of the selection. We report in particular two deviations from this optimal selection. First, we found a bias towards the options explored (Fig. 5): participants selected a higher proportion of explored options (68%, CI = [62.29%, 72.92%]) than an optimal agent would have done (64%, CI = [58.55%, 68.76%];  $W = 891.50$ ,  $p = .006$ ;  $r = .46$ , 95% CI = [.17, .67]).

Second, we found that participants occasionally selected an option that is numerically dominated by another one that was not selected. In other words, for the options they explored, they did not always select options with the highest total scores; for the options they did not explore, they did not always select the ones with the highest score A. Overall, these errors significantly reduced participants' performance by 1.27 percentage point (95% CI = [.54%, 2.00%];  $W = 521.00$ ,  $p < .001$ ;  $r = .65$ , 95% CI = [.43, .80]).

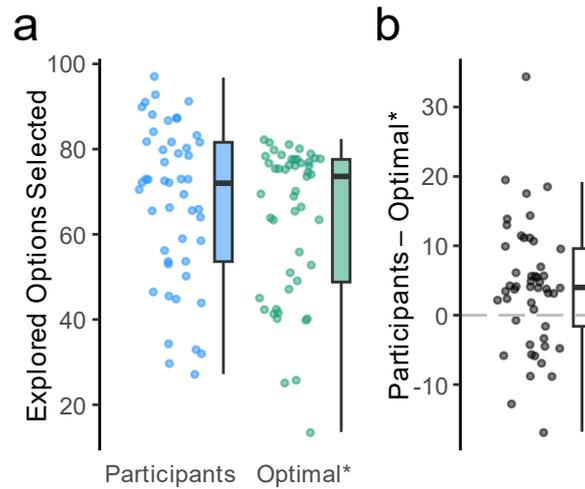


Fig 5.

Selection bias toward explored options (i.e. options with a revealed score  $B$ ). As we did for exploration bias (Fig. 4), we represented first (a) the proportions of explored options selected by participants and by the optimal agent considering participant exploration, and secondly (b) the difference of proportions between participant and this optimal agent selection, showing that participants over-selected the explored options.

## IV. Experiment 2

### a. Motivation

To examine the robustness of the bias towards exploring the highest options, in subsequent experiments we reduced the exploration set (Exp. 2A) and increased the selection set (Exp. 2B). Both manipulations make this bias more costly: optimal exploration moves away from the highest options, and exploring the highest becomes almost useless in terms of performance (Fig. 6). As a result, the normalized performance of the *Highest* strategy is significantly lower in Exp. 2A (79%:  $W = 705.00$ ,  $p = .001$ ;  $r = -.39$ , 95% CI  $[-.57, -.18]$ ) and 2B (50%:  $W = 7.50$ ,  $p < .001$ ;  $r = -.99$ , 95% CI  $[-1.00, -.99]$ ) compared to Experiment 1 (84%). Besides, as they enforce a clear distinction between exploration and selection sets, these manipulations eliminate the potential confusion for participants between the two stages. Moreover, as these manipulations make the task more challenging (as less information is available in Exp. 2A) or less challenging (as there is less room for errors in Exp. 2B), they allowed us to test the presence of this bias under various performance levels.

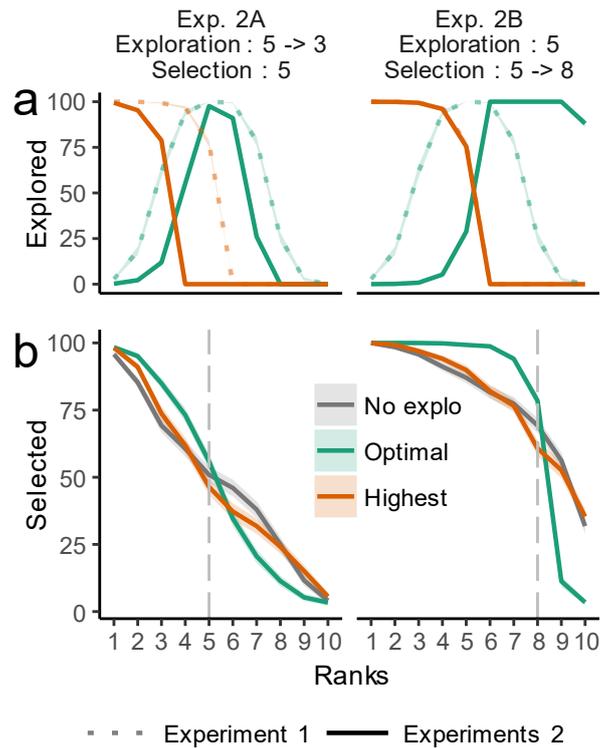


Fig 6.

Optimal exploration and selection in Experiment 2. Left panels correspond to Exp. 2A, where the number of explored options was changed from 5 to 3, and right panels correspond to Exp. 2B, where the number of options to select was changed from 5 to 8. The top row shows (a) the proportion of options explored by rank of score A for the Optimal and Highest strategies. Note that the Highest strategy exploration is identical whether the selection is 5 or 8. The bottom row (b) illustrates the proportions of options selected as a function of the final rank, under the Optimal exploration, Highest exploration, and No exploration strategies.

## b. Method

Based on the analyses conducted with G\*Power (Faul et al., 2007), 17 participants were needed to detect an effect size of 0.65 (corresponding to the effect size for the exploration bias in Experiment 1, see Fig. 4) with a statistical power of 80% and a false positive rate of 5%. We however recruited a larger sample, similar to that of Experiment 1, to facilitate the comparison of results between studies.

Experiment 2A involved 47 participants (24 women) and Experiment 2B involved 48 participants (27 women). Participants' mean age was 30 years (ranging from 18 to 67). In Experiment 2A, the payment scheme was as in Experiment 1, and the mean gain was 3.87€. In Experiment 2B, since it was impossible to give less than three correct responses, each correct response was rewarded 0.50€ instead of 1€, for a mean gain of 3.73€.

## c. Results

### Performance

As in Experiment 1, participants' performance was well above chance (2A: 77% vs. chance-level 54%:  $W = 1128.00$ ,  $p < .001$ ;  $r = 1.00$ , 95% CI = [1.00, 1.00] ; 2B: 93% vs. chance level at 83%:  $W = 1166.00$ ,  $p < .001$ ;  $r = .98$ , 95% CI = [.97, .99]) but below optimal performance (2A: 86% ,  $W = 981.00$ ,  $p < .001$ ;  $r = 0.98$ , 95% CI = [.97, .99]; 2B: 98%;  $W = 935.50$ ,  $p < .001$ ;  $r = .98$ , 95% CI = [.96, .99]). In comparison to an agent with no exploration phase, we found that participants did not exhibit better performance in Exp. 2A (77%,  $W = 617.50$ ,  $p = .403$ ), but they did in Exp. 2B (91%,  $W = 811.00$ ,  $p < .001$ ;  $r = .57$ , 95% CI = [.31, .75]).

Decomposing the loss in performance relative to the optimal agent revealed a significant loss both at the exploration stage (Exp. 2A: 4.39 percentage points,  $W = 934.00$ ,  $p < .001$ ;  $r = .97$ , 95% CI = [.95, .99]; Exp. 2B: 3.68 percentage points,  $W = 924.00$ ,  $p < .001$ ;  $r = .95$ , 95% CI = [.91, .98]) and at the selection stage (Exp. 2A: 4.90 percentage points,  $W = 789.00$ ,  $p < .001$ ;  $r = .83$ , 95% CI [.70, .91]; Exp. 2B: 1.55 percentage points,  $W = 738.00$ ,  $p < .001$ ;  $r = .63$ , 95% CI = [.40, .79]).

### Exploration stage

Having to explore 3 options instead of 5 in Exp. 2A made the optimal strategy more concentrated around the middle profiles (Fig. 6 and 7a1). Yet, participants still showed a significant bias towards exploring the highest profiles, which was observed on the aggregated data (Fig. 7a1), but also at the individual level when comparing participants' exploration to the optimal exploration ( $W = 841.00$ ,  $p = .003$ ;  $r = .49$ , 95% CI = [.21, .70], see Fig. 7b1). We noted however that this bias was less pronounced than in the previous experiment: whereas 65% of participants were classified as using the *Highest* heuristic in Exp. 1, this proportion fell to 40% in Exp. 2A ( $\chi^2 = 5.01$ ,  $p = .025$ ; Adjusted Cramer's  $v = .23$ , 95% CI = [.00, 1.00]). Participants who used the *Highest* heuristic consistently completed the exploration stage faster than other participants ( $W = 147.00$ ,  $p = .009$ ;  $r = -.45$ , 95% CI = [-.68, -.14]). In sum, changing the exploration set to 3 options lead to a reduction of the bias, as expected, but not a full elimination of it.

In Exp. 2B, having to pick 8 options out of 10 in the selection phase shifted the optimal exploration towards the lowest ranks (Fig. 6 and 7a2), and a clear shift was also observed when considering aggregated behaviour of participants (Fig. 7a2). Although the bias towards the highest profiles was still significant at the group level ( $W = 948.50$ ,  $p < .001$ ;  $r = .75$ , 95%

CI [.58, .86]), a closer examination revealed important inter-individual differences in this dataset, with many participants exhibiting little or no bias (see Fig. 7b2). Here, a majority (54%) of participants was in fact best described by the *Lowest* heuristic, which was close to optimal in this setting. The *Highest* heuristic was adopted by 33% of participants, a proportion that was non-negligible, but significantly below the 65% observed in Exp. 1 ( $\chi^2 = 8.68$ ,  $p = .003$ ; Adjusted Cramer's  $v = .30$ , 95% CI = [.11, 1.00]). These results confirmed that our manipulation largely affected participants' exploration. Participants using *Highest* strategy were marginally faster in their exploration than others ( $W = 174.00$ ,  $p = .074$ ;  $r = -.32$ , 95% CI [-.59, .02]).

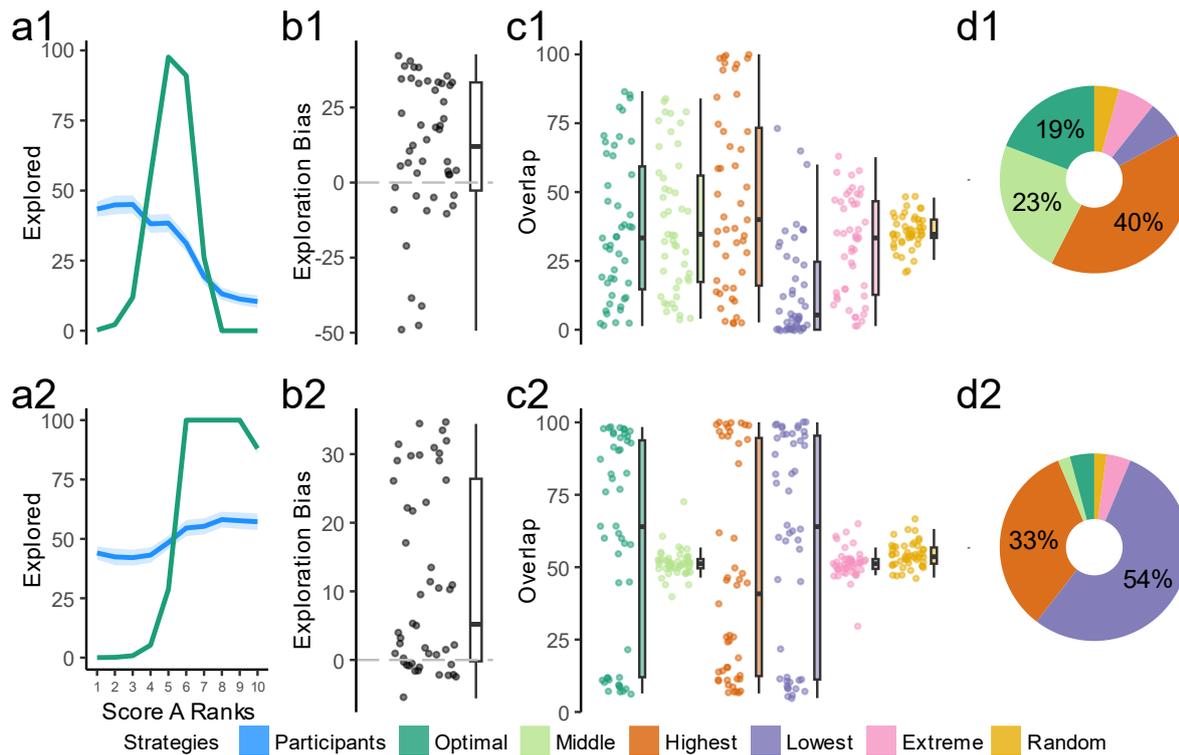


Fig. 7.

Participants' exploration compared with the optimal strategy and various heuristics in Experiment 2A (3 options explored, top row) and 2B (8 options selected, bottom row). This figure displays the proportions of explored profiles within each rank (from 1 to 10), for participants and for the optimal agent in Exp. 2A (a1) and Exp. 2B (a2), highlighting the difference of proportions between participant and optimal agent exploration in Exp. 2A (b1) and Exp. 2B (b2). We also represented the proportions of overlap between participants exploration and optimal exploration / heuristics in Exp. 2A (c1) and Exp. 2B (c2) explained by the proportions of the best-matching strategies, across participants in Exp. 2A (d1) and Exp. 2B (d2).

### Selection stage

As was found for Exp. 1, participants in both Exp. 2A and 2B selected a higher proportion of explored options than an optimal agent would have done (2A:  $W = 630.00$ ,  $p = .010$ ;  $r = .46$ ,  $95\% \text{ CI} = [.17, .68]$ ; 2B:  $W = 957.00$ ,  $p < .001$ ;  $r = .77$ ,  $95\% \text{ CI} = [.60, .87]$ ). In other words, we found again a bias towards the explored options during the selection phase.

For completeness, we also evaluated the prevalence of mathematical errors, by which participants selected an option numerically dominated by another one that was not selected. We found that such errors reduced participants' performance by 2.76 percentage points in Exp. 2A ( $W = 761.00$ ,  $p < .001$ ;  $r = .77$ ,  $95\% \text{ CI} = [.60, .87]$ ) and by .78 percentage point in Exp. 2B ( $W = 436.50$ ,  $p = .001$ ;  $r = .65$ ,  $95\% \text{ CI} = [.43, .80]$ ).

## V. Experiment 3

### a. Motivation

In Experiments 1 and 2 we documented a bias towards the highest options when exploring a set of candidates in the context of a set selection task. The goal of Experiment 3 was to distinguish between two interpretations of this observation: is exploration attracted towards the highest scores, or towards the options that will require an action in the selection stage? To disentangle these two possibilities, we changed the procedure of the selection stage and asked participants to act not on the options that they want to keep, but on the ones they want to reject. If exploration is biased towards the highest scores, then the sign of the bias should remain the same under this new procedure. If, however, the bias is towards options that require an action, then one should observe a reversal of the exploration bias.

### b. Method

Experiment 3 involved 52 participants (27 women), with a mean age of 35.5 years (ranging from 20 to 69).

In each trial, options within the rejection set that aligned with the worst set of 5 options were considered correct responses. To maintain a similar earning system, the options were ranked from last to first instead of from the first option to the last option as in Experiments 1 and 2. Thus, ties remained favourable to the participants, and there were always at least 5 options to reject. The overall earnings of participants were contingent on the performance of a randomly selected trial, with a mean gain of 4.13€.

The bias towards the highest options was calculated as before, such that a positive bias indicates a tendency to over-explore the highest options relative the optimal exploration (as was observed in previous experiments). A negative bias would correspond to over-exploration of the lowest options, indicating that the bias is aligned with the action required at the selection stage.

### c. Results

#### Performance

As in Exp. 1, participants' performance was greater than chance (79% vs. 54%:  $W = 1366.00$ ,  $p < .001$ ;  $r = .98$ , 95% CI = [.97, .99]). Participants tend to perform slightly better than an

optimal agent deprived of exploration (78%:  $W = 864.50$ ,  $p = .060$ ;  $r = .30$ , 95% CI = [.00, .56]), and they still do not reach optimal performance (92%:  $W = 1326.00$ ,  $p < .001$ ;  $r = 1.00$ , 95% CI = [1.00, 1.00]).

When decomposing performance loss, we found a significant loss both at the exploration stage (4.78 percentage points,  $W = 1302.00$ ,  $p < .001$ ;  $r = .96$ , 95% CI [.93, .98]) and at the rejection stage (8.03 percentage points,  $W = 1152.50$ ,  $p < .001$ ;  $r = .96$ , 95% CI = [.93, .98]).

### Exploration stage

Overall, participants in Experiment 3 have explored more low-ranked options than high-ranked options (Fig. 8a). However, when computing exploration biases for each participant, we found no significant effect at the group level ( $W = 637.50$ ,  $p = .642$ ; Fig. 8b), but a large inter-individual variability (see Fig. 8c and 8d). Indeed, most participants were best described by two opposite strategies: *Lowest* or *Highest*. Many participants still followed the *Highest* but this proportion has significantly reduced from 65% in Exp. 1 to 23% in Exp. 3 ( $\chi^2 = 16.62$ ,  $p < .001$ ; Adjusted Cramer's  $v = .42$ , 95% CI = [.24, 1.00]). At the same time, the proportion of participants using *Lowest* increased from 12% in Exp.1 to 44% in Exp.3 ( $\chi^2 = 11.10$ ,  $p < .001$ ; Adjusted Cramer's  $v = .34$ , 95% CI = [.16, 1.00]).

In contrast to this variability across individuals, further analyses indicated a strong stability of the best-matching strategy within each individual. Specifically, when the best-matching strategies were estimated for odd and even trials separately, we found that they were identical between the two subsets of trials for 69% of participants, largely exceeding chance level (95% CI = [23%, 41%], estimated with 10k permutations across participants).

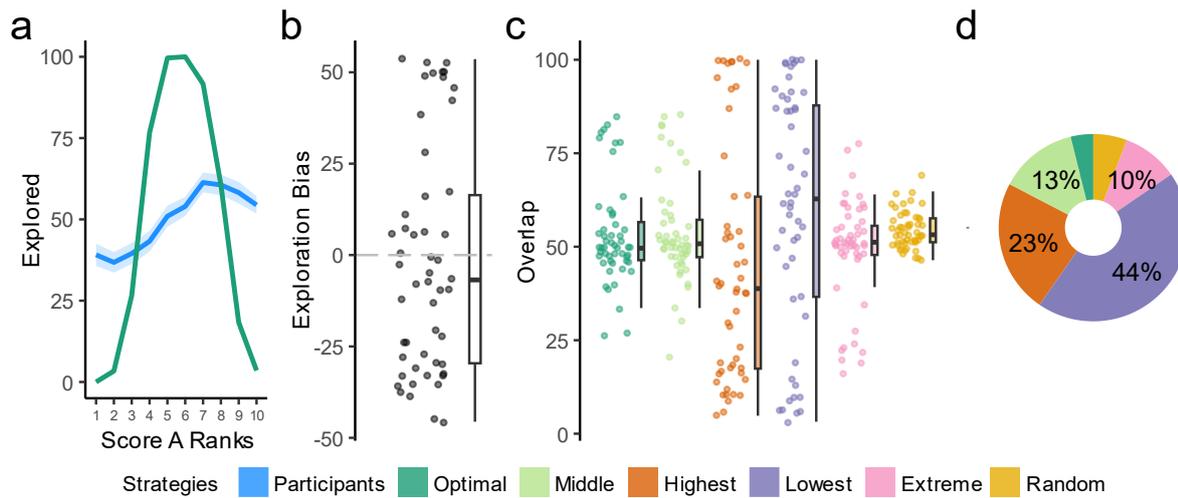


Fig. 8.

Participants' exploration compared with the optimal strategy and various heuristics in Experiment 3. As we did for Exp. 1 and 2, we represented (a) the proportions of explored profiles within each score A rank (from 1 to 10), for participants and for the optimal agent in Exp. 3, highlighting the elimination of (b) the difference of proportions between participant and optimal agent exploration of highest options. We also display (c) the proportions of overlap between participants exploration and optimal exploration / heuristics also represented as (d) the proportions of the best-matching strategies, across participants.

### Rejection stage

Although Exp. 3 was formally equivalent to Exp. 1, our decomposition of performance loss indicated the loss at the rejection stage became larger in Exp. 3 compared to Exp. 1 (8.03% vs. 4.64%). We found that a large part of this additional loss was due to more frequent mathematical errors (Exp. 3: 3.68% vs. 1.27% in Exp. 1;  $W = 1576.00$ ,  $r = .24$ ,  $p = .039$ ).

Regarding the bias towards the explored options, we found that participants in Exp. 3 rejected a lower proportion of explored options than an optimal agent would have done ( $W = 321.00$ ,  $p = .006$ ;  $r = -.45$ , 95% CI =  $[-.67, -.18]$ ). In other words, across the two different framings of the task (i.e. selection in Exp. 1 vs. rejection in Exp. 3), this bias corresponds to giving lower scores to unexplored options, leading to fewer selections (in Exp. 1) or more rejections (in Exp. 3) of these options.

## VI. Discussion

We designed a new two-stage task: Exploration of a set of options followed by selection of the best options based on available information. We demonstrated that participants' exploration exhibits a bias towards options with the highest scores. This bias is reduced when it becomes more costly in terms of performance, for example when fewer options must be explored or more options must be selected. The bias is even eliminated (at the group level) when the selection stage becomes a rejection stage: most participants then explore options with the lowest scores, although some still explore options with the highest scores. During the second stage (selection/rejection), a selection bias towards explored options is observed. This bias remains stable across the different variations tested.

The experimental literature on information search has focused on situations where a single option must be selected. At the theoretical level, the distinction between the two is more important than it seems: exploring the best options can be optimal for a single selection task but suboptimal in the case of set selection. At the practical level, we believe that the set selection task also corresponds to a range of situations that are relevant ecologically. A student may select several campuses to visit; a recruiter may select several candidates to invite for an interview, etc.

Exploring the favourite options seems to be most common in real life. Our task was inspired by the search for information in the domain of recruitment. In this domain, collecting information is crucial but difficult, and the information available is uncertain, which leaves ample room for the deep influence of stereotypes and prior beliefs. This influence has been documented in laboratory tasks and field studies, on measures such as decisions to invite a candidate for an interview, wages or promotion offers (Kroll et al., 2021; Neumark, 2018; Pager et al., 2009).

It has also been shown that stereotypes influence the exploration phase. For instance, in 3 field studies, Bartoš et al. (2016) show that individuals in a minority group received different amounts of attention compared to individuals in a majority group. They describe two types of scenarios. In the labour market where recruiters are selective regarding the number of applications they want to consider, they pay more attention to applications from the majority group. In a rental market where landlords are relatively less selective regarding the number of applicants they want to invite to their apartment, they scrutinize with more attention the applications coming from the minority group. These results resonate with our experimental findings in Experiment 1 and in Experiment 2B, respectively.

Prior research on discrimination in labour markets sought to identify whether it results from preferences towards one group or another, or from a rational statistical discrimination (Guryan & Charles, 2013; Phelps, 1972). In Bartoš et al. (2016), the authors propose that their results are consistent with a rational decision process, in which agents would have different prior beliefs regarding the quality of applications from the majority and minority groups. Given our findings, however, we may suggest an alternative interpretation where statistical discrimination is not a rational exploration strategy but results from a bias towards the highest scores.

One important issue for future research is to better understand the computational mechanisms of this bias towards the favourite. One could speculate that participants have a favourable belief towards the options with a high score, and that a pavlovian bias (Hunt et al., 2016), a confirmation bias or a positive testing strategy (see e.g. Klayman 1995) may lead them to explore preferentially these options. Another explanation could be that people are attracted to non-instrumental information to increase their confidence in their decision (Eliaz & Schotter, 2007, 2010; Matthews et al., 2023).

We note that although most participants failed to adopt the optimal exploration strategy in our paradigm, some participants were able to find a close approximation, which was to explore the middle ranks in Experiments 1, 2A and 3 (and the lowest ranks in Experiment 2B). Thus, they achieve a "satisfactory" performance without having to pay the high computational cost of determining exactly the optimal choices (Gigerenzer & Gaissmaier, 2011; Lee & Cummins, 2004; Petit et al., 2021). Understanding how these participants reasoned would be useful to help others improve their exploration strategy. They may have realized that seeking additional information for the top and bottom candidates might be useless because these candidates will most likely end up in the top 5 and in the bottom 5 anyway and concluded that the middle strategy is a good exploration strategy.

A key question for future research would be to help participants adopt such reasoning in the laboratory or in the field. Alternatively, one could engineer decision aids that directly guide participants towards exploring the options with the highest instrumental value of information. We note that such recommendations may differ from recommendations based on the most popular or favourite contents, typically used in commercial websites or streaming platforms.

The ethical acceptability is also a crucial consideration in the design of recommendations. In our paradigm, it was possible for participants to explore the option initially ranked 6<sup>th</sup>, and if they then realized that this option would perform well, to select this option over, say, an option initially ranked 3<sup>rd</sup> which was not explored. In a real recruitment process, this would be problematic. If a candidate X is a priori better than candidate Y, it would be unacceptable to only invite candidate Y for the interview, and to hire Y but not X, after observing that Y performed very well during this interview. In that situation, candidate X would complain that they were not given a fair chance because they were not invited for an interview. In our paradigm, one could conjecture that this ethical constraint might not change much the optimal exploration strategy. Whether it may change participants' exploration behaviour remains an empirical question.

## Bibliography

1. Bartoš, V., Bauer, M., Chytilová, J., & Matějka, F. (2016). Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition. *American Economic Review*, 106(6), 1437–1475. <https://doi.org/10.1257/aer.20140571>
2. Bouhlel, I., Chessa, M., Festré, A., & Guerci, E. (2022). When to stop searching in a highly uncertain world? A theoretical and experimental investigation of “two-way” sequential search tasks. *Journal of Economic Behavior & Organization*, 203, 80–92. <https://doi.org/10.1016/j.jebo.2022.08.034>
3. Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97. <https://doi.org/10.1016/j.jbef.2015.12.001>
4. Eliaz, K., & Schotter, A. (2007). Experimental Testing of Intrinsic Preferences for Noninstrumental Information. *The American Economic Review*, 97(2), 166–169.
5. Eliaz, K., & Schotter, A. (2010). Paying for confidence: An experimental study of the demand for non-instrumental information. *Games and Economic Behavior*, 70(2), 304–324. <https://doi.org/10.1016/j.geb.2010.01.006>
6. Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G\*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/BF03193146>
7. Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic Decision Making. *Annual Review of Psychology*, 62(1), 451–482. <https://doi.org/10.1146/annurev-psych-120709-145346>
8. Guryan, J., & Charles, K. K. (2013). Taste-Based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots. *The Economic Journal*, 123(572), F417–F432. <https://doi.org/10.1111/eoj.12080>
9. Hausmann, D., & Läge, D. (2008). Sequential evidence accumulation in decision making: The individual desired level of confidence can explain the extent of information acquisition. *Judgment and Decision Making*, 3(3), 229–243. <https://doi.org/10.1017/S1930297500002436>
10. Hunt, L. T., Rutledge, R. B., Malalasekera, W. M. N., Kennerley, S. W., & Dolan, R. J. (2016). Approach-Induced Biases in Human Information Sampling. *PLOS Biology*, 14(11), e2000638. <https://doi.org/10.1371/journal.pbio.2000638>
11. Juni, M. Z., Gureckis, T. M., & Maloney, L. T. (2016). Information sampling behavior with explicit sampling costs. *Decision*, 3(3), 147–168. <https://doi.org/10.1037/dec0000045>
12. Kassambara, A. (2023). *Ggpubr: 'ggplot2' Based Publication Ready Plots. R package version 0.6.0.* <https://CRAN.R-project.org/package=ggpubr>
13. Klayman, J. (1995). Varieties of Confirmation Bias. In J. Busemeyer, R. Hastie, & D. L. Medin (Eds.), *Psychology of Learning and Motivation* (Vol. 32, pp. 385–418). Academic Press. [https://doi.org/10.1016/S0079-7421\(08\)60315-1](https://doi.org/10.1016/S0079-7421(08)60315-1)

14. Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature Human Behaviour*, 3(6), 587–595. <https://doi.org/10.1038/s41562-019-0589-3>
15. Kroll, E., Veit, S., & Ziegler, M. (2021). The Discriminatory Potential of Modern Recruitment Trends—A Mixed-Method Study From Germany. *Frontiers in Psychology*, 12. <https://www.frontiersin.org/articles/10.3389/fpsyg.2021.634376>
16. Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the “take the best” and the “rational” models. *Psychonomic Bulletin & Review*, 11(2), 343–352. <https://doi.org/10.3758/BF03196581>
17. Lüdecke, D., S. Ben-Shachar, M., Patil, I., M. Wiernik, B., Bacher, E., Thériault, R., & Makowski, D. (2022). *easystats: Framework for Easy Statistical Modeling, Visualization, and Reporting*. CRAN. <https://easystats.github.io/easystats/>
18. Matthews, J. R., Cooper, P. S., Bode, S., & Chong, T. T.-J. (2023). The availability of non-instrumental information increases risky decision-making. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-023-02279-1>
19. Neumark, D. (2018). Experimental Research on Labor Market Discrimination. *Journal of Economic Literature*, 56(3), 799–866. <https://doi.org/10.1257/jel.20161309>
20. Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of “success”. *Journal of Behavioral Decision Making*, 17(2), 117–137. <https://doi.org/10.1002/bdm.465>
21. O. Wilke, C. (2024). *Cowplot: Streamlined Plot Theme and Plot Annotations for 'ggplot2'*. R package version 1.1.3. <https://CRAN.R-project.org/package=cowplot>
22. Pager, D., Bonikowski, B., & Western, B. (2009). Discrimination in a Low-Wage Labor Market: A Field Experiment. *American Sociological Review*, 74(5), 777–799. <https://doi.org/10.1177/000312240907400505>
23. Petitet, P., Attaallah, B., Manohar, S. G., & Husain, M. (2021). The computational cost of active information sampling before decision-making under uncertainty. *Nature Human Behaviour*, 5(7), 935–946. <https://doi.org/10.1038/s41562-021-01116-6>
24. Phelps, E. S. (1972). The Statistical Theory of Racism and Sexism. *The American Economic Review*, 62(4), 659–661.
25. R Core Team. (2023). *R: A Language and Environment for Statistical Computing*. <https://www.R-project.org/>
26. Rassin, E., Eerland, A., & Kuijpers, I. (2010). Let’s find the evidence: An analogue study of confirmation bias in criminal investigations. *Journal of Investigative Psychology and Offender Profiling*, 7(3), 231–246. <https://doi.org/10.1002/jip.126>

27. Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99. <https://doi.org/10.2307/1884852>
28. Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., ... Yutani, H. (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, 4(43), 1686. <https://doi.org/10.21105/joss.01686>
29. Zeileis, A., Meyer, D., & Hornik, K. (2007). Residual-Based Shadings for Visualizing (Conditional) Independence. *Journal of Computational and Graphical Statistics*, 16(3), 507–525. <https://doi.org/10.1198/106186007X237856>

## Supplementary materials

### Expected performance, optimal selection and optimal exploration

For both participants' responses and for the different strategies introduced above, we computed expected performance as follows.

#### **Selection stage**

In a given trial, the actual performance of a given selection ranges between 0% and 100% and is defined as the number of selected options that indeed belong to the best options in that trial, divided by the number of selections (i.e. 5). The expected performance of a selection is defined as the expectation of this performance over all possible cases compatible with the observed scores in that trial (i.e. over all  $11^5$  possible values for unrevealed scores B).

#### *Optimal selection*

The optimal selection set is defined as the selection set that maximizes the expected performance, for a given exploration. We computed this optimal selection and its associated expected performance for each trial.

#### **Exploration stage**

We defined the expected performance of an exploration set as the performance that would be obtained at the selection stage, if this exploration was followed by an optimal selection. Note that in theory it is defined in expectation over all possible values of scores B, i.e.  $11^{10}$  possibilities. In practice, determining the expected performance for a given exploration set is computationally too intensive and impractical for the entirety of our experiments. Therefore, to estimate this variable, we considered not all possible realisations of scores B but only a fraction. Specifically, we sampled .5% of the possible realisations for the revealed scores B (i.e. 805 realizations), and for those realisations we defined the optimal selection by considering all  $11^5$  possible realisations for the remaining (i.e. unrevealed) scores B. Simulations indicated that when estimating expected performance in this way, the variability in the estimated value was less than .05%. Note that we evaluated several exploration strategies in our analyses using this approach, and we used the same realizations (within a given trial) across the different strategies to evaluate their expected performance.

*Optimal exploration*

Through tests on reduced cases (4 profiles and scores between 0 and 4) which we could explore exhaustively, we found that optimal explorations are always a set of options with consecutive ranks on score A, and more specifically the options closest to the average of the 5<sup>th</sup> and 6<sup>th</sup> values on score A (or 8<sup>th</sup> and 9<sup>th</sup> values for Experiment 3 in which 8 options must be the selected). When searching for the optimal exploration strategy in our main data, for each trial we thus tested only explorations that verified this property. Among these, the optimal exploration strategy was defined as the one that maximises the expected performance.

Strategies overlaps

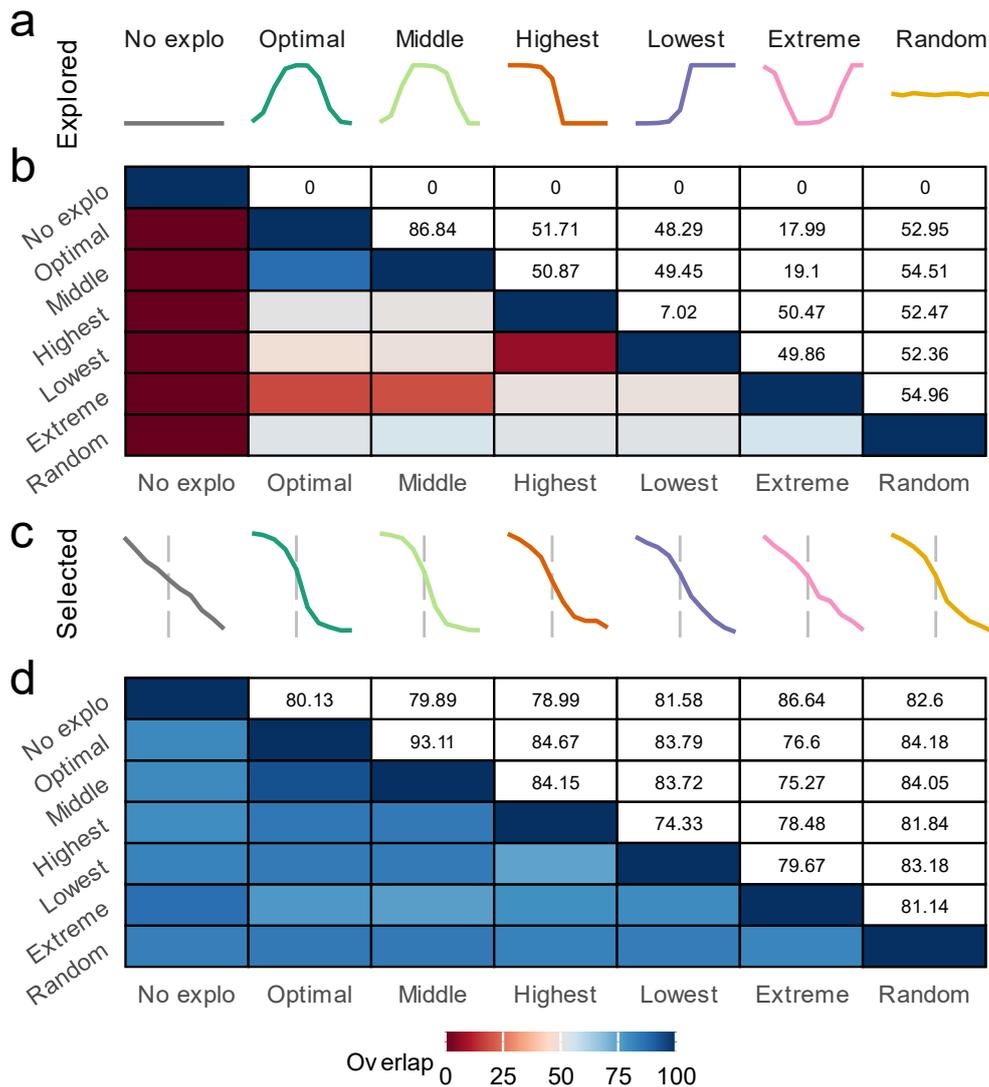


Fig.9.

Description of heuristics and optimal strategies and overlaps between them at exploration and selection stages for Experiment 1. These graphics displays (a) the proportion (from 0 to 100%) of explored options per scores A rank (from 1 to 10) for heuristics and optimal strategies, (b) the proportion of explored options overlapping between strategies during the exploration stage, (c) the proportion (from 0 to 100%) of selected options per final rank (from 1 to 10) for heuristics and optimal strategies, and (d) the proportion of selected options overlapping between strategies during the selection stage.

Exploration heuristics with correlated scores A and B

The optimal strategy presented here would consistently outperform any other heuristic if scores A and B were correlated. Figure 10 shows the performance improvement for each strategy compared to an agent who had not explored at all. Selection is conducted by taking into account the explored B scores and the given correlation. Notably, we see that the Highest strategy is not more optimal when the scores are correlated.

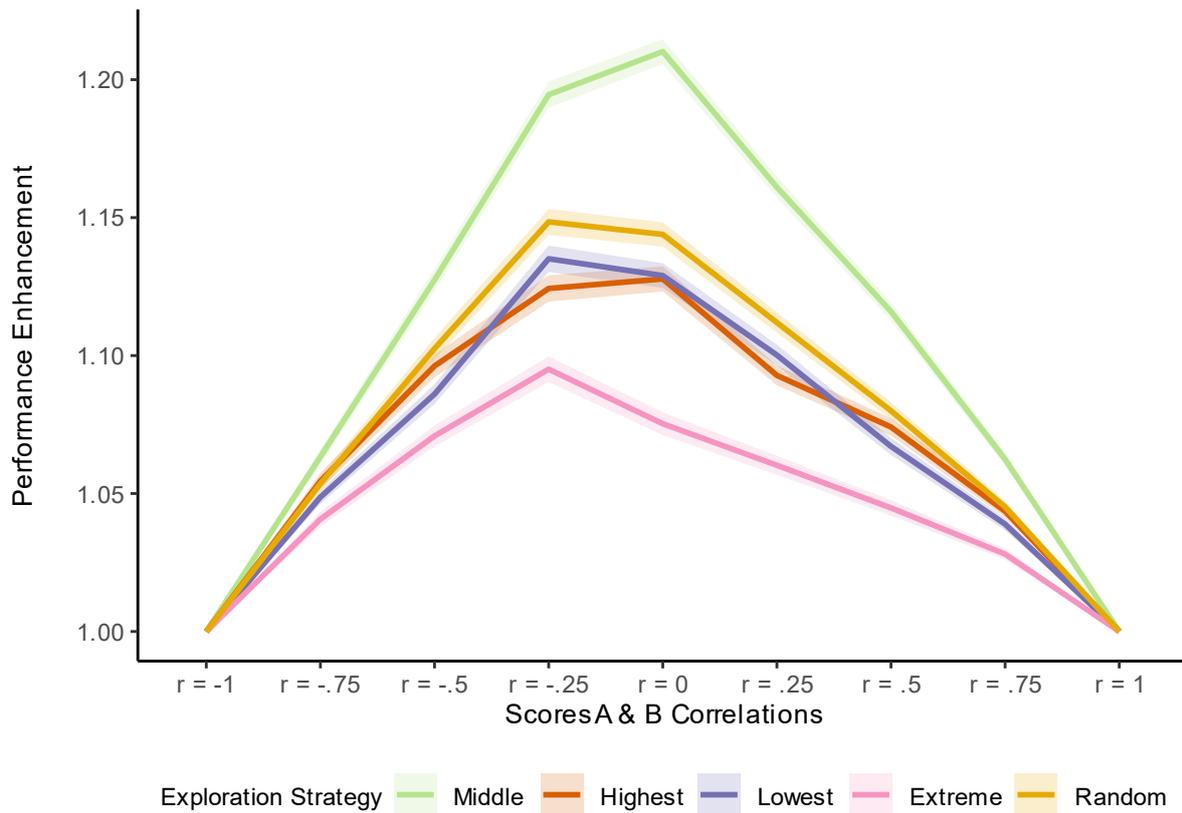


Fig 10.

Graph representing the performance of exploration heuristics as a function of the correlation between scores A and B. Performance is shown as the ratio between the performance achievable with a given exploration strategy and the performance achievable by an agent with no exploration. The correlations between scores A and B are generated with a probability ( $p$ ) that score B is equal to score A (e.g.,  $-.25 \rightarrow p = .25$ ), and a probability ( $1-p$ ) that score B is an integer drawn from a random distribution between 0 and 10. For negative correlations,  $pp$  is the probability that score B is symmetrical to score A relative to 5 (e.g., score A = 7  $\rightarrow$  score B = 3).

# Chapter 2: Cognitive Factors Influencing Information Search Biases

Jean-Michel Dagba<sup>1,2</sup>, Alexandre Lietard<sup>3,4</sup>, Jean-Christophe Vergnaud<sup>1,5</sup>, Vincent de Gardelle<sup>1,5,6</sup>

<sup>1</sup> Centre d'Économie de la Sorbonne (CES) & Université Paris 1 Panthéon Sorbonne, France

<sup>2</sup> Entreprise Humans Matter, France

<sup>3</sup> Université Paris Cité, France

<sup>4</sup> KU Leuven, Belgium

<sup>5</sup> Centre national de la recherche scientifique (CNRS), France

<sup>6</sup> Paris School of Economics (PSE), France

## Table of contents

Table of contents	64
Introduction	65
Do Set Selection Paradigms Reveal Distinct Cognitive Biases?	65
Cognitive biases in experts' decision-making	68
Experiment 4: Biases and heuristics tasks	69
Method	69
Results	74
Summary of experiment 4	80
Experiment 5: Expertise	81
Methods	81
Results	82
Summary of experiment 5	83
General Discussion	84
Summary of the Study and Main Findings	84
Interpretation of the Results	86
Limitations of the Study	87
Future Directions and Conclusion	87
Bibliography	89
Supplementary materials	92
The Questionnaire Version Replicates the Original Task Version	92

## I. Introduction

The primary goal of this chapter is to understand individual differences in decision-making within our set selection paradigm. Specifically, we aim to identify factors that contribute to exploration bias, selection bias, and selection errors observed in the task. By doing so, we can determine whether these phenomena align with established mechanisms in the literature or if they represent novel findings. To achieve this, we conducted two experiments that examine the role of cognitive biases and expertise in decision-making.

Furthermore, individual differences in the expression of these biases suggest that there is no uniform response across participants, as seen in our previous findings. To better understand these dynamics, we draw an experiment on categorisations of biases and test relevant factors, to assess how they contribute to decision-making in our experimental paradigm (Exp. 4).

The literature has also examined expertise as a source of individual differences in the expression of cognitive biases. We designed a second experiment to measure the effect of expertise on the expression of the exploration and selection biases identified in our paradigm (Exp. 5).

### a. Do Set Selection Paradigms Reveal Distinct Cognitive Biases?

#### **Confirmation bias**

One particular focus of our study is the potential relationship between exploration bias towards favourite options and *confirmation bias*. Intuitively, the exploration bias towards favourites seems similar to a form of confirmation bias, which refers to the tendency of individuals to seek information that confirms their hypothesis (Klayman, 1995; Klayman & Ha, 1987; Wason, 1960, 1968).

This is one of the most well-known biases that affects information search. In Wason's (1968) classic task demonstrating this bias, participants must seek information to verify a rule of logical implication ( $P \rightarrow Q$ ). They show a tendency to explore information that could confirm the proposed logical rule (checking that Q occurs when P occurs), but not information that could disprove it (checking that P does not occur when non-Q occurs).

In our task, it may be that participants explore the candidates with the highest to confirm that they are indeed in the top 5. However, it is true that our task may be functioning differently because it does not ask to test a logical implication (such as  $P \rightarrow Q$ ). Those, it remains an empirical question to test if our task involves confirmation bias.

Confirmation bias can be measured through various well-known tasks in the literature. We use three of these in our experiment (4-Cards Selection Task, Interviewee's Personality Task, and the 2-4-6 Task). These tasks allow us to measure the tendency to confirm a hypothesis rather than to disprove it.

### **Framing effect**

*Framing effect* is another cognitive bias that could be related to the exploration bias towards favourites. The framing effect is the tendency to be influenced by the way information is presented (Kahneman & Tversky, 1984). As shown by the results of Exp. 3 in Chapter 1, when the objective is presented as eliminating the 5 worst candidates, the majority exploration strategy is to explore the candidates with the lowest scores A. A significant proportion of participants still choose to explore the highest scores A, as does the majority when the objective is framed as finding the top 5. It therefore seems relevant to investigate whether there is a link between the exploration bias towards favourites and the framing bias, as known in the literature.

### **Risk aversion**

Finally, one might wonder whether the selection bias towards explored options is a form of *risk aversion*. Since participants tend to slightly favour candidates for which they know both scores, rather than only the A score, this could indicate a perception of lower risk with explored candidates. An explored candidate has at least a known ranking among all explored candidates. Its final ranking can only be modified by the non-explored options. For the unexplored options, however, uncertainty is complete, as the score could still increase by up to 10 points. It is possible for an unexplored option to finish first or last without knowledge of its B score. Risk aversion might thus manifest as a selection bias towards explored options, while risk-seeking behaviour would be reflected by a selection bias towards unexplored options. We included a lottery task, created by (Holt & Laury, 2002), to measure risk aversion.

### **Cognitive Biases Factors**

For practical reasons, we selected a set of bias tasks to test in our experiment, based on the relevant biases identified earlier (confirmation, framing, risk aversion) and the pertinent factors presented in studies categorising cognitive biases. Many recent studies have listed and categorised cognitive biases by investigating their common factors and their links with one (or more) explanatory factor related to decision-making abilities (Aczel et al., 2015; Arnott, 2006; Berthet & de Gardelle, 2023; Ceschi et al., 2019; Dimara et al., 2020). These studies provide two main insights relevant to our research: First, there are individual

differences in the expression of cognitive biases (Aczel et al., 2015; Berthet, 2021; Stanovich et al., 2008; Teovanović et al., 2015), just as there are in cognitive abilities. This echoes our findings from Chapter 1, where the majority of participants employed a clearly biased exploration strategy in one direction, others adopted a bias in the opposite direction, and a minority of participants exhibited no bias at all. Second, there is no consensus on a categorisation of biases that would clearly identify common mechanisms. As described by Ceschi et al. (2019), different approaches—each with their own strengths and weaknesses—have been proposed to create a taxonomy of cognitive biases (cognitive, empirical, theoretical approaches, etc.). Thus, different approaches lead to various categorisations.

### *Some relevant factors to explore*

For our study, we chose to draw primarily from the categories proposed by Stanovich et al. (2008), Ceschi et al. (2019) and Berthet, Autissier, & De Gardelle, (2022), which introduce factors particularly relevant to our task and are coherent with each other. Their grouping into factors suggests links between tasks, so we decided to retain two tasks per factor to verify whether they fall under the same factor. We selected biases with well-established behavioural tasks in the literature, with multiple items and good reliability (Berthet & de Gardelle, 2023).

***Mindware gaps.*** The mindware gaps factor is described as a lack of knowledge (or ability) to apply probabilistic rules and relevant knowledge (Ceschi et al., 2019). This factor seems relevant to explore as our task is based on calculating probabilities from rules that must be understood to solve the task. This factor encompasses seven biases in the factor analysis reported by Ceschi et al. (2019). The concept is also referenced in Stanovich et al.'s (2008) inventory as “Probability Knowledge.” We kept *base-rate neglect*, the tendency to overemphasise specific information at the expense of base rates (Kahneman & Tversky, 1973), and the *conjunction fallacy*, i.e. the tendency to believe that a conjunction of events is more likely than the events comprising it (Tversky & Kahneman, 1983). These biases could be related to participants' ability to understand and correctly apply the probabilistic rules presented in the instructions (the independence of scores and sampling from a uniform distribution without replacement).

***Anchoring and adjustment.*** The anchoring & adjustment factor is defined as the tendency to be influenced by a reference point (Ceschi et al., 2019). It closely resembles the “Focal Bias” factor described by Stanovich et al. (2008). This factor groups four tasks in Ceschi et al.'s (2019) factor analysis, including two well-known tasks from the literature: *anchoring bias* and the *framing effect*.

Since there is no clear and unified categorisation of confirmation bias tasks and we used three different tasks to assess this bias, we did not add a specific task to associate with confirmation bias.

Finally, we included the *Cognitive Reflection Test (CRT)* and the *Berlin Numeracy Test (BNT)*. Cognitive biases are generally weakly correlated with cognitive ability tasks according to the literature (Berthet, 2021; Stanovich & West, 2008; Teovanović et al., 2015). To explore whether cognitive abilities are a distinct factor, we added the Cognitive Reflection Test (CRT), a series of questions designed to assess the ability to avoid an intuitive (incorrect) response in favour of an easy-to-calculate correct response. Similarly, we added the Berlin Numeracy Test (BNT), which is also a well-established cognitive ability task in the literature. The BNT designed by Cokely et al. (2012) assesses statistical numeracy and risk comprehension.

### b. Cognitive biases in experts' decision-making

A factor that can influence our cognition and behaviour is the expertise acquired in a specific domain. Expertise in a field allows individuals both to reduce errors and to be more optimal through learned strategies. Recent literature has begun to examine cognitive biases in specific professional populations, such as medicine, finance, or law (Blumenthal-Barby & Krieger, 2015; Hodgkinson et al., 1999; Kumar & Goyal, 2015; Lidén et al., 2019). Overall, research shows that experts are not immune to cognitive biases, and they are all overconfident (Berthet, 2022). These studies are relatively recent, and it remains to be determined whether an expert's performance in a bias task is linked to the expression of that same bias in their professional context.

This experimental framework was designed to create an original recruitment task that captures the complexity of sourcing, a practice that serves as the starting point for many recruitment processes. Studies show that sourcing on the web to find potential candidates to a job, particularly on LinkedIn, is rapidly expanding (Abbas et al., 2021). This raises many questions about how people are screened and selected from the vast array of relevant and accessible profiles online (D'Silva, 2020; Roulin & Levashina, 2019).

To our knowledge, the scientific literature does not yet tell us whether recruiters are also influenced by cognitive biases in their work, as other individuals are. Although a few studies have pointed to suboptimal CV reading processes (Cole et al., 2007), most recruitment research focuses more on the social aspect (such as gender or racial discrimination) rather

than the cognitive aspect. In this study, we aim to address the first step, which is to measure the biases of expert participants in information search and set selection, namely recruiters.

## II. Experiment 4: Biases and heuristics tasks

### a. Method

#### Study overview

Figure 1 displays the order of task administration. The experiment required approximately 50 minutes to complete (including 15 min for the main task and around 2-3 min for each subsequent task) and was remunerated based on performance from a randomly chosen trial in the main task + the outcome of a randomly chosen lottery.

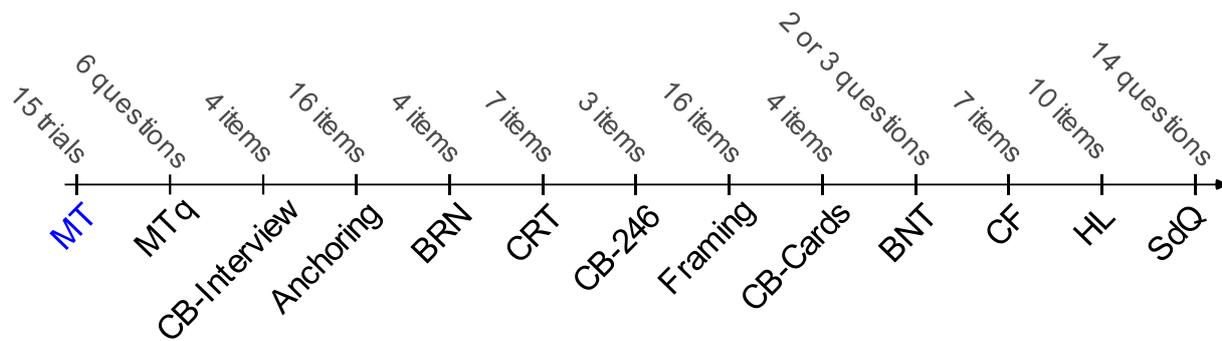


Fig. 1.

Illustration of the timeline for Experiment 4. MT = Main task of set selection (~15 min to complete), MTq = Main task in a form of single choice questionnaire, CB-Interview = Confirmation bias Interviewee's Personality Task, Anchoring = Anchoring bias task, BRN = Base Rate Neglect task, CRT = Cognitive Reflective Test, CB-246 = Confirmation bias 2-4-6 task, Framing = Framing effect task, CB-Cards = Confirmation bias 4 cards-selection task, BNT = Berlin Numeracy Test, CF = Conjunction Fallacy task, HL = Holt-Laurry lotteries task, SdQ = Sociodemographic questionnaire.

#### Main task (MT)

The main task used was the main task presented in chapter 1. To recap, participants face a recruitment task where they initially receive one piece of information (score A) on the 10 profiles presented, and then a second piece of information (score B) for 5 profiles of their choice among the 10 (exploration stage). They must then choose the 5 profiles with the highest total score (score A + score B) from the 10 profiles (selection stage). Participants are informed that scores A and B are equally important, independent, and drawn from a uniform distribution between 0 and 10. The unexplored B scores are just as important as the explored scores B in determining their performance on a trial (see Methods in Chapter 1 for details).

Participants are required to complete three comprehension questions regarding the instructions, followed by one training trial and then 15 trials. The questions during the instructions serve to ensure that the participants have acquired an understanding of the underlying probability rules of the task, namely the independence of scores and the sampling with replacement from a uniform random distribution. The questions are presented in the form of a single choice among three possible responses. In the event of an incorrect response, the participants are provided with an explanation and must provide the correct answer on subsequent trials until they do so.

After the instructions, participants complete a training trial during which they receive feedback on the number of selected profiles that are among the best, the revelation of the scores B they did not explore, and the complete ranking of the 10 profiles. This ensures that participants understand how their performance is defined. During the remainder of the experiment, the feedback after each trial contains only the number of selected profiles that are among the best profiles.

### Main task measures

The measures and models employed in Chapter 1 are reused:

**Performance:** The percentage of options chosen by the participant that are actually ranked between the 1st and 5th positions.

**Exploration strategies:** The categorisation of exploration strategies is based on the same heuristics (*Highest, Lowest, Middle, Extreme, Random*) and the same method of estimating the Optimal exploration strategy.

**Exploration bias:** The difference in the proportion of options ranked among the best after score A that were explored by the participants compared to the same proportion explored according to the Optimal. A positive value indicates a bias in exploration towards the favourite options.

**Selection bias:** The difference between the proportion of explored options selected by the participants and the same proportion selected according to an optimal selection following the same exploration. A positive value indicates a selection bias towards the explored options.

**Mathematical errors:** The percentage of chosen options that were dominated by another equally explored option among the selected options.

### Main Task in a form of simple choice questionnaire (MTq)

After completing 15 trials of the main task, participants complete a theoretical version of the same task, without any financial incentives for performance. Details and analysis of this task are reported in the Supplementary materials.

### *Confirmation biases (Interview, Cards, 246)*

Next, the participants completed confirmation bias tasks inspired by Wason's original work (1960, 1968) and adapted into a multiple-choice questionnaire in French by Berthet et al. (2022). These tasks assess individuals' ability to test hypotheses. In the 2-4-6 task (CB-246), participants are required to find the rule that generates a sequence of three numbers. They are initially provided with a sequence that conforms to the rule (2, 4, 6), and they must determine if the rule is "add 2." To do this, they are presented with two new sequences of numbers: one that confirms the suggested rule (8, 10, 12) and one that disconfirms it (3, 6, 9). For each of these sequences, they can decide whether to propose it or not (Yes or No) to determine if the rule "add 2" is correct. The score is the difference between the response to the confirming sequence (Yes = 1 and No = 0) and the disconfirming sequence. The task score is the average of the scores from the three items.

The *Four-Card Selection Task (CB-Cards)* involves presenting participants with a rule of logical implication of the form  $P \rightarrow Q$  (e.g., "If a card has a D on one side, then it has a 5 on the other side"). Participants have four elements: P, not-P, Q, not-Q, which they can choose to flip or not (cards D, 7, 5, and K). To verify if the rule is true or false, it is necessary and sufficient to turn over the P and not-Q elements among the four (in this case, the D and 7 cards). The score for an item is the difference between the number of Q and not-Q elements flipped (either 1, 0, or -1). The task score is the average of the scores from the four items. For both tasks, a positive score indicates a tendency to confirm hypotheses (confirmation bias), and a negative score indicates a tendency to disconfirm hypotheses.

The last is the *"Interviewee's Personality Task" (CB-Interview)* designed by Snyder & Swann (1978) and adapted by Berthet et al. (2022). Participants were given a hypothesis about a job candidate's personality (e.g., "the candidate is introverted"). To test the hypothesis, participants had to choose 4 questions from a list of 10 questions (4 confirming the hypothesis, 4 disconfirming it, and 2 neutral). The score for an item is the number of questions chosen that confirm the hypothesis. The total score is the average of the scores from the four items. The higher the score, the greater the confirmation bias.

In all three experiments, significant confirmation biases were found, with p-values  $< .001$  and  $r > .71$  for all tasks (Wilcoxon signed rank test with continuity correction). None of the three tasks included financial incentives based on performance.

### *Anchoring bias*

We utilised the version of the task by Jacowitz & Kahneman (1995) adapted by Berthet, et al. (2022), without financial incentives for performance. The measures and the 16 items are identical to those used by Berthet et al. (2022). The task involves (1) estimating whether a

real-world quantity is greater or smaller than an anchor (for example, “Has Norway won more or fewer than 25 Nobel Prizes?”), and then (2) providing a free estimate. Since the responses are open-ended, there are no theoretical bounds. To address outliers, values beyond  $\pm 1.5$  IQR are replaced by the value of the nearest quartile. Among the 83 participants, 5 had a score below the expected range between 0 and 1. We also found a significantly positive anchoring effect greater than 0 ( $W = 3455.00$ ,  $p < .001$ ;  $r = 0.98$ ), indicating that participants' numerical estimates were influenced by the anchor introduced in the question.

#### *Base-Rate Neglect (BRN)*

We utilized the task and measures from Erceg et al. (2022), an adaptation of the task originally designed by De Neys & Glumicic (2008), without financial incentives for performance. The task consists of 4 items, each presenting a description followed by a binary choice between a correct response option (statistically more probable) and an incorrect response option (statistically less probable). Participants achieved an average performance of 64% correct responses, slightly higher than participants in Erceg et al.'s study (mean = .64,  $SD = .38$  vs. mean = .37,  $SD = .40$ ).

#### *Cognitive Reflection Test (CRT)*

We translated into French the 3 questions proposed by Frederick (2005) and added 4 new questions. The task involves answering open-ended questions with a single correct answer, a numerical value. A participant's score is the sum of their correct answers. Participants were not given any financial incentive based on the quality of their answers in this task. The average performance is 53%, which is highest to the 41% performance reported by Frederick (2005). In Table 1, we report the scores obtained for our extended version of 7 questions, whose average performances are also fairly close (40% in Experiment 1 and 47%).

#### *Framing effect*

We utilized the task and measures from Berthet (2021) based on Bruine De Bruin et al. (2007), without financial incentives for performance. The task consists of 8 pairs of dilemmas with opposite framings (gains vs. losses), where participants respond on a 6-point Likert scale ranging from 1 (“I choose option A without hesitation,” the risky option) to 6 (“I choose option B without hesitation,” the certain option). As predicted by Berthet (2021) based on prospect theory, we found a significantly positive average score ( $W = 1877.00$ ,  $p < .001$ ;  $r = 0.80$ ), indicating a tendency to opt more for the risky option in loss frames and the certain option in gain frames.

### *Berlin Numeracy Test (BNT)*

It is an adaptive test that contains a total of 4 open-ended questions, each requiring a unique exact numerical answer. Participants answer question 1, then proceed to question 2b if their response is correct, and to question 2a if it is incorrect. An incorrect response to question 2b leads to question 3. Participants score 1 if they answer question 2a incorrectly, a score of 2 if they answer question 2a correctly, a score of 3 if they answer question 3 incorrectly, and a score of 4 if they answer question 2b or question 3 correctly. The average scores are 2.06 (SD = 1.18) and 2.54 (SD = 1.17) in Exp. 1 and 4, respectively, which is consistent with the average score of 2.6 (SD = 1.13) reported by Cokely et al. (2012). This task did not contain any financial incentives based on performance.

### *Conjunction fallacy (CF)*

We used the 7 items from the task (translated into French) and measures from Burgoyne et al. (2023a), without financial incentives for performance. Participants read a scenario and estimate which of two options (event A OR a conjunction of events A and B) is more likely to occur. Participants performed 1.6 times better than those in Burgoyne et al.'s study (mean = .53, SD = 2.16 vs. mean = .33, SD = 1.93).

### *Loteries (HL)*

The next task is a series of 10 pairs of lotteries, designed by Holt & Laury (2002). Each pair of lotteries consists of a low-stake lottery (win of €1.60 or €2.00) and a high-stake lottery (win of €0.10 or €3.85). The highest gain (€2.00 and €3.85) has a probability  $P$  of being won, and the lowest gain has a probability of  $1-P$ .  $P$  takes the values 0.1, 0.2, ... 1. For each of the 10 pairs of lotteries, participants must choose to bet on either the low-stake lottery or the high-stake lottery. Participants receive the payout from one of the ten lotteries they chose, randomly selected. For the analyses, the number of high-stake lotteries chosen will be recorded. The higher this value, the lower the risk aversion.

### *Socio-demographic Questionnaire*

The socio-demographic questionnaire allows us to collect information such as age, gender, level of education (1 = High school diploma or less, 2 = Some college or equivalent, 3 = Bachelor's degree, 4 = Beyond Bachelor's degree), and participants' confidence in their understanding of probabilities and mathematics (from 1 = not confident to 5 = confident). We also gathered professional information to identify any potential expertise in recruitment, as the main task is presented as a recruitment task.

## Participants

Experiment 4 involved 83 participants (44 women). Participants' mean age was 26.3 years (ranging from 18 to 64). The mean payment was 15.40€ including a participation fee of 9€.

## Apparatus

Experiments were conducted online. They were build using oTree (Chen et al., 2016) with the technical support of the Fédération S2CH.

## Statistical analysis

Statistical analyses were conducted using R (R Core Team, 2023), as well as the packages 'tidyverse' (Wickham et al., 2019), 'ggpubr' (Kassambara, 2023), 'cowplot'(O. Wilke, 2024), 'easystats' (Lüdtke et al., 2022), 'jtools' (Long, 2024) and 'vcd' (Zeileis et al., 2007).

## b. Results

### Instructions checks

Overall, 80% of the participants made no errors on the 3 comprehension questions regarding the probability rules used in the experiment. Only 16% of participants made 1 error, and 4% made 2 errors. The performance of participants who made no errors is not significantly better than that of participants who made 1 or 2 errors ( $W = 623.50$ ,  $p = .484$ ;  $r = 0.11$ , 95% CI = [- .20, .40]). This indicates that the experiment was well understood, including the underlying probability rules.

### Replication

First, we replicated all the main results of the Chapter 1 Exp. 1, as detailed below.

#### *Performance*

The mean participants performance in Experiment 4 was 84%, which is less than the optimal performance of 91% ( $W = 3207.00$ ,  $p < .001$ ;  $r = .98$ , 95% CI = [.97, .99]), and not different to the participants performance in Experiment 1 ( $W = 1749.50$ ,  $p = 0.181$ ;  $r = -.14$ , 95% CI = [- .33, .06]).

### *Exploration strategies*

In Experiment 4, as in Experiment 1, a large majority of participants used a strategy close to Highest (73% in Exp. 4, 65% in Exp. 1). Participants using this strategy were significantly faster to complete the exploration stage ( $W = 185.00$ ,  $p < .001$ ;  $r = -.72$ , 95% CI =  $[-.83, -.56]$ ), suggesting that this strategy is the simplest or most intuitive to use. In contrast, only a minority of participants were closer to the Optimal exploration strategy (8% in Exp. 4, 12% in Exp. 1). The distribution of participants according to their preferred exploration strategy did not significantly differ between Experiments 1 and 4 (Fisher's Exact Test for Count Data,  $p$ -value = 0.2488).

### *Exploration bias*

Like Experiment 1, participants in Experiment 4 significantly explore more of the favourite options (ranked between 1 and 5 according to their score A) than would be expected by the optimal strategy ( $W = 3081.00$ ,  $p < .001$ ;  $r = 0.81$ , 95% CI =  $[.71, .88]$ ). The exploration bias is similar to that observed in Experiment 1 ( $W = 2268.50$ ,  $p = .269$ ;  $r = 0.12$ , 95% CI =  $[-.09, .31]$ ).

### *Selection bias*

As expected, participants also exhibit a significant selection bias ( $W = 2389.50$ ,  $p < .001$ ;  $r = .48$ , 95% CI =  $[.26, .64]$ ). They tend to select options they have explored more than they should, similar to Experiment 1 ( $W = 1898.00$ ,  $p = .525$ ;  $r = -.07$ , 95% CI =  $[-.26, .14]$ ).

### *Mathematical errors*

Finally, participants make mathematical errors that significantly reduce their performance by 1 percentage point ( $W = 1052.00$ ,  $p = .008$ ;  $r = .42$ , 95% CI =  $[.19, .60]$ ), which is equivalent to the loss observed in Experiment 1 ( $W = 1839.00$ ,  $p = .352$ ;  $r = -.10$ , 95% CI =  $[-.29, .11]$ ).

## **Factors influencing exploration bias, selection bias and maths errors in the task**

Through our analyses, we highlight relationships between maths errors in our task and scores on the Cognitive Reflection Test, the Berlin Numeracy Test, and the Conjunction Fallacy task. These three tasks appear to cluster under a single factor.

Our analyses focus on the 10 tasks: Cognitive Reflection Test (*CRT*), Confirmation Bias (*CB-Cards*, *CB-246*, *CB-Interview*), Risk-aversion Lotteries (*HL*), Berlin Numeracy Test (*BNT*), *Anchoring*, *Framing*, Base Rate Neglect (*BRN*), and Conjunction Fallacy (*CF*). As presented in Tab. 1, all tasks showed good reliability scores (Cronbach's alpha), except for the anchoring

task, with average scores relatively consistent with findings from previous research (see details in the Methods section). According to normality tests, no task has a score that follows a normal distribution (as indicated by the Skewness and Kurtosis values in Tab. 1). Through correlations, exploratory factor analysis and regressions, we aim to determine whether the cognitive abilities engaged in these tasks are also involved in our paradigm.

Task	Experiment 4 (N=83)								
	<i>items</i>	<i>range</i>	<i>mean</i>	<i>SD</i>	<i>median</i>	<i>MAD</i>	<i>SKW</i>	<i>Kurt.</i>	<i>reliability</i>
CRT	7	[0, 7]	3.30	1.87	4.00	1.48	-.30	-.88	/
BNT	2 or 3	[1, 4]	2.54	1.17	2.00	1.48	-.01	-1.48	/
CB-246	3	[-1, 1]	.57	.55	.67	.49	-1.32	4.02	.77
CB-Cards	4	[-1, 1]	.43	.58	.50	.74	-.82	2.77	.82
CB-Interview	4	[0, 4]	2.02	.70	2.00	.74	.28	2.90	.66
Anchoring	16	NA	.31	.22	.26	.22	.46	2.41	.32
BRN	4	[0, 1]	.64	.38	.75	.37	-.60	1.88	.83
Framing	8	[-5, 5]	.73	1.12	.25	.56	1.65	6.13	.94
CF	7	[0, 7]	3.70	2.16	4.00	2.97	-.31	2.11	.76
HL	10	[0, 10]	4.43	1.94	4.00	1.48	.28	1.47	/

Tab. 1. Descriptive statistics of cognitive biases tasks. The CRT, BNT, and HL consist of a series of items that can be summed but are not designed to calculate an inter-item reliability score. For HL, we report only the number of decisions favouring the riskier option of the two lotteries presented. Range represents the interval of possible scores. The anchoring task involves open responses, so there is no theoretical range. Standard deviation (SD) represents the dispersion around the mean. Median absolute deviation (MAD) indicates the dispersion of values around the median. Skewness (SKW) measures the degree of asymmetry in the distribution relative to a normal distribution ( $> 0$  for right skew,  $< 0$  for left skew). Kurtosis (Kurt.) measures the concentration of values around the mean and the thickness of the distribution tails compared to a normal distribution ( $> 0$  for a sharper distribution,  $< 0$  for a flatter distribution). Reliability is given by Cronbach's alpha coefficient, where values closer to 1 indicate stronger internal consistency.

### Correlations between the Tasks

We established a correlation matrix between the tasks, presented in Fig. 2. The Benjamini-Hochberg method was applied to correct the p-values as the analyses are primarily exploratory. There is several significant correlations between the tasks. Notably, there is a strong negative correlation between the CRT, reflecting the ability to inhibit an intuitive response, and CF, the tendency to overestimate the probability of joint events ( $r = -.49$ , 95% CI =  $[-.64, -.30]$ ,  $t(81) = -5.03$ , adj.  $p < .001$ ).

Among the variables from the main task, exploration bias, selection bias and mathematical errors are considered variables of interest. There is no significant correlation between these three variables of interest. Maths errors negatively correlated with CRT ( $r = -.31$ , 95% CI  $[-.49, -.10]$ ,  $t(81) = 2.91$ , adj.  $p = .032$ ) and positively correlated to CF ( $r = .29$ , 95% CI =  $[.08,$

.48],  $t(81) = 2.77$ , adj.  $p = .040$ ). To further explore the links between tasks and examine the factors identified in the literature, we conducted an exploratory factor analysis.

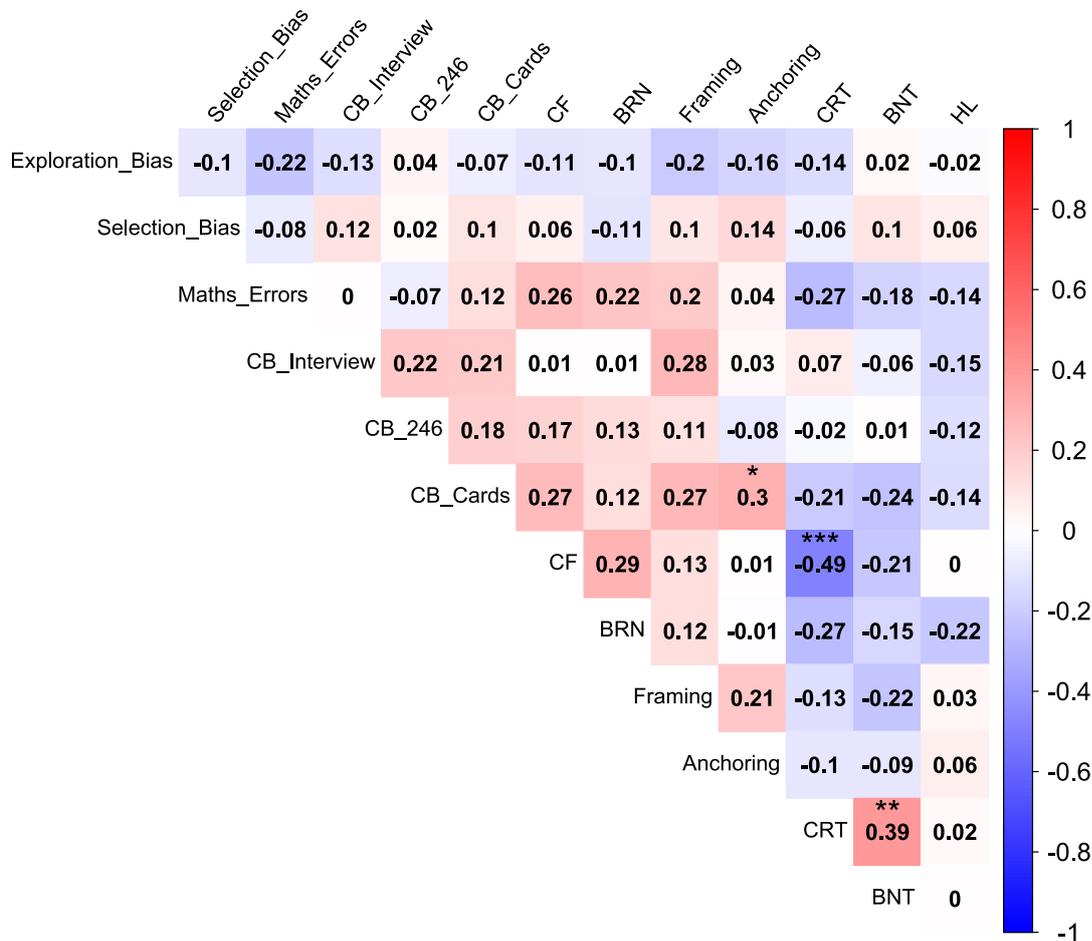


Fig. 2. Correlation matrix between the tasks used (CRT, BNT, CB-246, CB-Interview, CB-Cards, HL, BRN, CF, Anchoring, Framing), exploration bias, selection bias and maths errors in the set selection task. Significance levels after correction using the Benjamini-Hochberg method: < .05 \* < .01 \*\* < .001 \*\*\*

### Exploratory Factor Analysis (EFA)

The exploratory factor analysis allows us to identify coherent factors for performing regressions. According to preliminary tests, the data are suitable for conducting the EFA:  $KMO = .65$ ; Bartlett's test of sphericity is significant ( $Chisq = 119.5347$ ,  $Df = 55$ ,  $p < .001$ ). Since there are different valid methods for determining the appropriate number of factors for the data, we employed several to reach a consensus. Various evaluation methods indicated that the data fit best with a 3-factor model. Kaiser's criterion suggested a 3-factor model (3 eigenvalues greater than 1). Parallel analysis, which compares the real data with simulated data, also indicated 3 factors, as shown in Fig. 3a. Finally, the Sample-Adjusted Bayesian Information Criterion (SABIC), which helps to avoid model overfitting, also suggested a 3-

factor model (SABIC = -8.51). Graphically, we also observed an "elbow point" at 2 factors. Therefore, we explored both 2-factor and 3-factor models, which are presented in Figure 3.

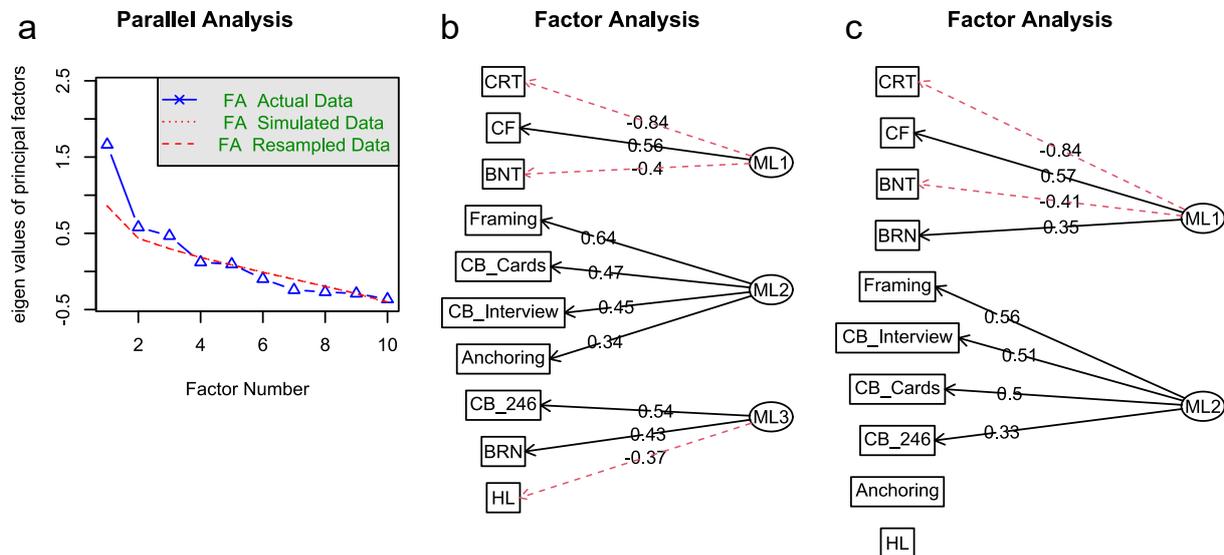


Fig. 3. Graphs presenting the factor analysis. Panel (a) represents the parallel analysis used to define the appropriate number of factors to retain in the model. Panel (b) presents the 3-factor model from the exploratory factor analysis. Panel (c) presents the 2-factor model.

In the 2-factor model, we distinguish one factor that is primarily influenced by confirmation bias and anchoring and framing tasks (ML2), and another factor that groups numerical and probabilistic reasoning tasks (ML1). These factors are consistent with the literature. Anchoring and framing tasks are grouped under an Anchoring & Adjustment factor and CF and BRN are grouped under a *Mindware Gaps* factor (Ceschi et al., 2019), *Confirmation Bias* tasks are grouped together, the CRT and BNT are close to each other (Cokely et al., 2012). This model explains 26% of the variance (Mean item complexity = 1.3, TLI = .944, RMSEA index = .028, BIC = -86.94). The 3-factor model explains 33% of the variance (Mean item complexity = 1.5, TLI = 1.176, RMSEA index = 0, BIC = -65.76). The first factor is loaded by CRT, BNT, and CF tasks, the second factor by Anchoring, Framing, CB-Cards, and CB-Interview tasks. The third factor is loaded by HL, CB-246, and BRN tasks.

### Linear Regression Models

Both regression models (estimated using OLS) significantly explain a moderate proportion of the variance in *Maths Errors* but not in *Exploration Bias* and *Selection Bias*, as illustrated by Fig. 4.

The 3-factor linear model predicts a significant proportion of the variance in *Maths Errors* ( $F(3, 79) = 4.36, p = .007, \text{adj. } R^2 = .11$ ). In particular, the first factor has a significant positive

effect (beta = 1.27, 95% CI = [.49, 2.04],  $t(79) = 3.26$ ,  $p = .002$ ; Std. beta = .34, 95% CI = [.13, .55]). The other two factors have no significant effect. This model is not significant for exploration bias ( $F(3, 79) = .91$ ,  $p = .442$ ), nor for selection bias ( $F(3, 79) = 0.46$ ,  $p = .713$ ). No factor has a significant effect for selection and exploration biases.

To detect potential effects of the factors, we compare the BIC of models by integrating the factors one by one. For the exploration bias, the intercept-only model (BIC = 753) has a better fit than the models integrating one of the three factors (BIC = [756, 757]) and models integrating combinations of two of the three factors (BIC = [760, 761]). The three-factor model (BIC = 763) most likely has a worse fit than the intercept-only model. For the selection bias, the intercept-only model also has a better fit (BIC = 598) than the model integrating one of the three factors (BIC = 602 for each) or a combination of two of the three factors (BIC = 606 for each). The model integrating all three factors has the worst fit (BIC = 610). It is therefore unlikely that the factors in this three-factor model allow for a better understanding of the variance in exploration bias or selection bias. For Maths Errors, the best fit is found for the model integrating only the first factor (BIC = 433) compared to other models integrating one, two, or three factors (BIC = [436, 446]).

The 2-factor model also significantly predicts the variance in *Maths Errors* ( $F(2, 80) = 6.43$ ,  $p = 0.003$ ,  $\text{adj. } R^2 = .12$ ), mainly due to the first factor, which could be assimilated to *Mindware Gaps* (beta = 1.32, 95% CI = [.55, 2.08],  $t(80) = 3.41$ ,  $p = .001$ ; Std. beta = .36, 95% CI = [.15, .56]). As we could see in the correlation matrix, *Maths Errors* seem particularly related to the CRT, BNT, and CF tasks, present in the first factor of each model, and BRN present in the first factor of the 2-factor model. There is no significant relation between the 2-factor model and *Exploration Bias* or *Selection Bias*.

We compare the BIC of models by integrating the factors one at a time. For exploration bias, the intercept-only model has a better fit (BIC = 753) than models integrating one factor (BIC = [755, 757]) or both factors (BIC = 759). For selection bias as well, the intercept-only model has a better fit (BIC = 598) than models integrating one of the two factors (BIC = [602, 603]) or both factors (BIC = 607). It is therefore unlikely that the factors in this two-factor model provide a better understanding of the variance in exploration bias and selection bias. For Maths Errors, the model integrating the first factor has a better fit (BIC = 432) than the intercept-only model (BIC = 439), the model integrating the second factor (BIC = 443), and the model integrating both factors (BIC = 436).

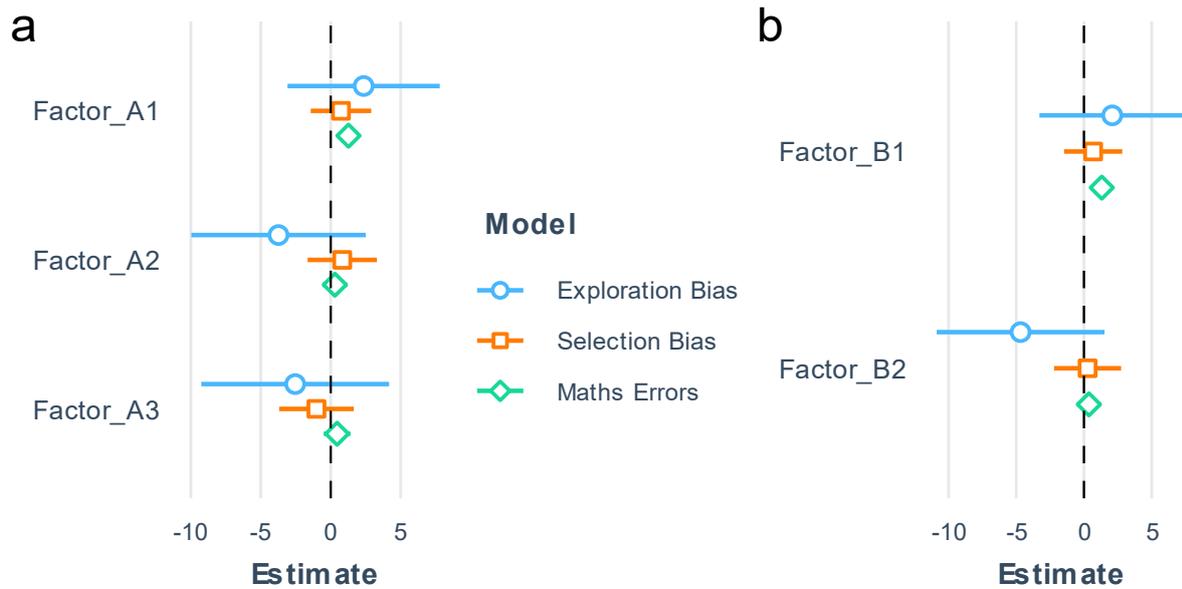


Fig. 4. Graphs of the regression models with (a) 3 factors and (b) 2 factors for exploration bias, selection bias, and maths errors in the set selection task.

### Summary of experiment 4

The aim of this experiment was to explore individual differences in the set selection paradigm by investigating cognitive factors. To achieve this, 83 participants completed a series of tasks (CRT, BNT, CB-Cards, CB-Interview, CB-246, Anchoring, Framing, BRN, CF, HL) after performing the set selection task. The main findings of this experiment are:

1. The regression models with cognitive factors explain a significant portion of the variance in maths errors.
2. However, exploration bias and selection bias do not appear to be linked to cognitive factors identified, either through the 2- or 3-factor regression models or through correlations with all tasks.

Experiment 5 below allows us to explore expertise as another potential factor influencing exploration bias.

### III. Experiment 5: Expertise

#### a. Methods

##### **Task**

Exp. 5 consisted solely of completing the questionnaire version of the original task (MTq). It is a theoretical version of the same task, without any financial incentives for performance. This made the experiment achievable in 5 to 10 minutes. This version of the task was tested in Experiment 4 and gave consistent results with the original version (more details in Supplementary materials). Since participants did not perform the original version of the task prior to the questionnaire, they received a theoretical presentation of the problem (available in Supplementary Materials) before choosing the exploration strategy they considered most optimal from among five heuristics (Highest, Lowest, Middle, Random, Extreme). The choices were presented in a random order to each participant. Participants then indicated their confidence in their response on a scale (1 = I guessed, 2 = I am not sure, 3 = I am fairly sure, 4 = I am certain). They were then asked for a second response: what would be the best strategy if the one they had initially chosen was not actually the best. They again indicated their confidence in this second response. Finally, participants completed a socio-demographic questionnaire and reported their level of expertise in recruitment.

##### **Participants**

The participants in Exp. 5 were recruited via LinkedIn and the participant pool of the Réseau Information sur les Sciences de la Cognition (RISC). Recruitment on LinkedIn was done through a public post aimed at recruitment professionals. The LinkedIn call for participants was shared by 11 people and viewed by over 15,000 individuals in total. This post primarily reached recruitment professionals. Recruitment via RISC helped supplement the pool with non-expert participants. In total, 169 individuals completed Exp. 5. Participants provided their consent to participate in an experiment of approximately 5 minutes, with no financial compensation.

##### *Expertise*

Recruitment expertise was self-reported by participants on a scale ranging from 1 (= Not at all expert) to 7 (= Completely expert). Participants who reported a level of 5, 6, or 7 were classified as experts (N = 78, or 46%). As shown by the histogram in Fig. 5, level 5 created a clear demarcation in the participant population.

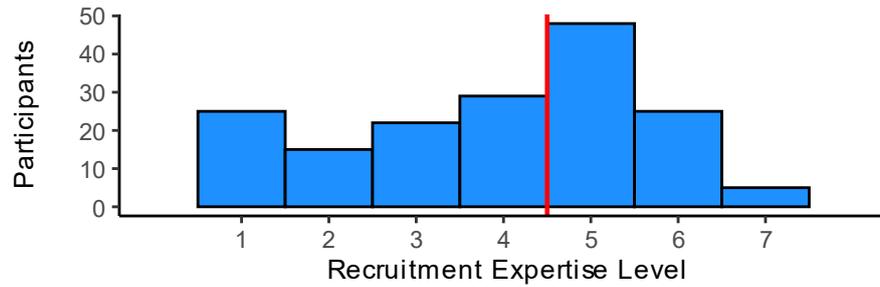


Fig. 5. Histogram showing the distribution of participants on the recruitment expertise scale.

## Apparatus

Experiments were conducted online. They were build using oTree (Chen et al., 2016) with the technical support of the Fédération S2CH.

## Statistical analysis

Statistical analyses were conducted using R (R Core Team, 2023), as well as the packages 'tidyverse' (Wickham et al., 2019), 'ggpubr' (Kassambara, 2023), 'cowplot'(O. Wilke, 2024), 'easystats' (Lüdecke et al., 2022), and 'vcd' (Zeileis et al., 2007).

## b. Results

### Recruitment Experts also Prefer the Heuristic Biased Towards Favourites

In Experiment 5, the questionnaire responses do not differ between the expert and non-expert groups (Fig. 6). For the first response, the distribution of heuristics chosen does not differ by group ( $\chi^2 = 1.63$ ,  $p = .803$ ). The most commonly chosen heuristic is Highest in both groups (experts: 33% vs. non-experts: 42%), consistent with all previous experiments. Similarly, the proportion of participants who find the correct heuristic among their two responses does not differ between groups ( $\chi^2 = .22$ ,  $p = .642$ ). The distribution of exploration strategy choices in the first response is significantly different from random ( $\chi^2 = 16.24$ ,  $p = .003$ ). However, the proportion of participants who found the correct heuristic among their two responses is not different from chance (chance level: 40% vs. 46%;  $\chi^2 = .54$ ,  $p = .462$ ).

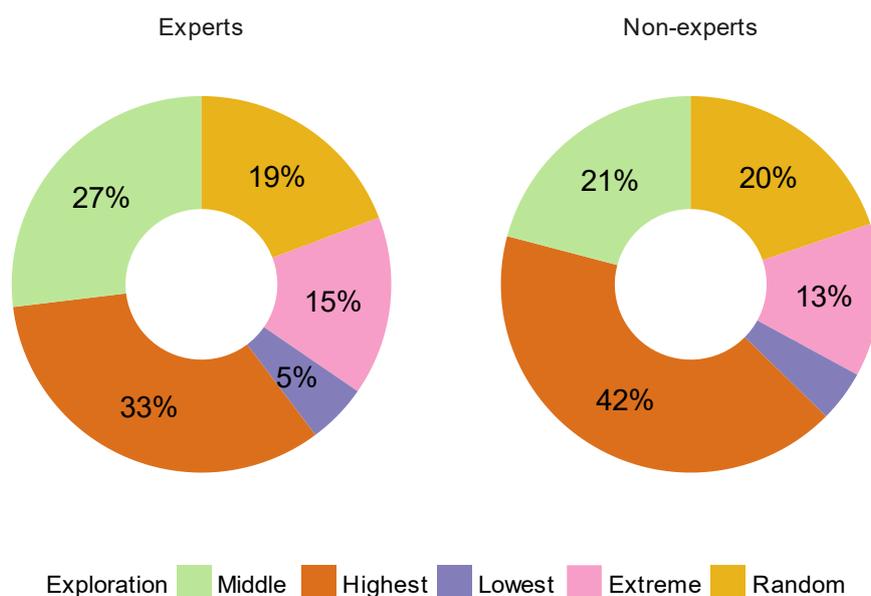


Fig. 6. Graphs representing exploration heuristic choices according to recruitment expertise in Experiment 5. The possible choices were Highest (explore candidates with the highest A scores), Lowest (explore the lowest), Middle (explore the middle), Extreme (explore both the highest and lowest), or Random (explore randomly).

### Confidence Does Not Indicate Expertise or Exploration Strategy

Reported confidence does not depend on the quality of the chosen strategy or recruitment expertise. The linear model indicates that the confidence reported after the first choice of strategy in the questionnaire is not explained by being self-assessed as an expert or having chosen the correct response ( $R^2 = 0$ ,  $F(2, 166) = .39$ ,  $p = .678$ ). The average confidence reported after the first response is significantly higher than the confidence reported after the second response (2.72 vs. 2.51,  $W = 1471.00$ ,  $p < .001$ ;  $r = .56$ , 95% CI = [.42, .66]). This large effect indicates that confidence judgments were likely not reported randomly. When the response time<sup>1</sup> for choosing a strategy increases, the confidence reported in this response decreases ( $r = -0.17$ ,  $p = .034$ ). The time taken to estimate confidence was not correlated with the confidence reported ( $r = .05$ ,  $p = .526$ ).

### Summary of experiment 5

Experiment 5 compares the exploration strategy choices of a population of recruitment experts (N=78) with a non-expert population (N=91). As individuals accustomed to searching

<sup>1</sup> For across-participants analysis, we excluded 13 participants who were outliers for the response time measure. Outliers are participants whose response time was outside the bounds (1.5 times the interquartile range above the third quartile or below the first quartile).

for profiles to create candidate shortlists, recruiters are assumed to be experts in the set selection paradigm. The main findings are:

1. Recruitment experts are just as likely to use a strategy biased towards favourites as non-experts.
2. An equal proportion of experts and non-experts find the best heuristic.
3. Confidence in the choice is not related to the quality of the response or to recruitment expertise.

The results of Experiments 4 and 5 are discussed in detail in the general discussion.

## IV. General Discussion

### a. Summary of the Study and Main Findings

This study aimed to investigate individual differences in decision-making within the set selection paradigm, particularly focusing on individual differences thought cognitive biases and recruitment expertise. To this end, two experiments were conducted, including a shortened version of the task that made it feasible for recruitment professionals, who are typically not part of lab-based participant pools. This adaptation allowed us to assess their exploration strategy in a controlled environment. We have five main findings in the Chapter, including three about the robustness of the set selection paradigm. First, we present the findings related to the robustness of the paradigm, followed by results linked to the individual factors tested.

The first key finding is the successful replication of performances and biases observed in Chapter 1. By reducing the number of trials from 25 to 15, participants in Exp. 4 (N=83) achieved performances similar to those seen in Experiment 1, confirming the robustness of the original findings. In both versions of the task, participants exhibited the same exploration and selection biases. Importantly, approximately 20-25% of participants adopted the optimal exploration strategy, while the majority preferred to explore their favourite options.

The second finding relates to the understanding of the task. Of the three questions verifying their understanding of probability rules, 80% of participants made no mistakes. Overall, participants understood the probability concepts used in the task well. Those who made some errors on these verification questions (and thus received feedback before proceeding) did not perform worse than others. This dismisses the idea that suboptimal performance is due to a lack of attention or a misunderstanding of the probability rules used.

The third finding about the robustness of the task concerned replication through the questionnaire version. It also enabled us to replicate results within participants (via Exp. 4) and across participants (via Exp. 5). When completing both versions of the task (in Exp. 4), participants generally selected the option in the questionnaire version that closely matched the strategy they had applied in the trials, and they were aware that their choice (in the questionnaire) matched their behaviour (in the original version). When completing only the questionnaire version without any previous experience with the task (Exp. 5), most participants chose the exploration strategy biased toward the favourites. The proportion of participants choosing the correct answer was equivalent to the proportion of participants exploring optimally in the original version.

The next results concern the study of factors that could explain individual differences. In Exp. 4, the scores on the various bias and cognitive ability tasks are consistent with the literature (Berthet & de Gardelle, 2023; Cokely et al., 2012; Frederick, 2005; Holt & Laury, 2002).

The fourth main finding is that Exploratory Factor Analysis identified factors that significantly explain individual differences in set selection task mathematical errors. With a 2-factor model, the regression could explain 14% of the variance in maths errors, thanks to the first factor loaded by CRT, BNT, CF, and BRN. With a 3-factor model, the regression model could explain 11% of the variance, thanks to the first factor loaded by CRT, BNT, and CF.

Exploration and selection biases in the set selection task were not correlated with any of the tested tasks related to cognitive abilities (CRT, BNT), risk aversion (HL), and cognitive biases (BRN, confirmation bias tasks, CF, Anchoring, Framing). Factor analyses did not yield a factor that significantly explained the variance in exploration and selection biases.

The last finding is that the study of recruitment expertise revealed that recruiters are just as prone to using a biased exploration strategy. Like non-experts, recruiters preferred to explore the options with the highest scores (i.e. the favourites), demonstrating the same exploration bias. We found individual differences in the exploration strategy chosen, but a relatively equal proportion of experts and non-experts chose the correct answer. Despite individual differences in strategy choices, confidence levels were moderate and not predictive of response quality or expertise.

## b. Interpretation of the Results

We demonstrated that the Chapter 1 results are robust and generalisable. In Experiment 4, we found the same levels of performance, exploration, and selection biases, after verifying task comprehension. The results of Chapter 1 were therefore not due to a misunderstanding of the task, which strengthens the findings of Chapter 1.

The questionnaire version used in Exp. 4 and Exp. 5 shows that individuals explicitly and consistently choose the exploration strategy biased towards the favourites, despite the presence of other response options, including the correct answer. In Exp. 4, where participants completed both versions of the task, they explicitly chose the same strategy in both formats, demonstrating that their behavioural strategies (in the original version) were explicit choices that they also selected in the theoretical version (the questionnaire). This persistence in exploring favourites indicates that participants may believe they are employing the best strategy, even when presented with more efficient alternatives.

The presence of the same bias among experts (recruiters) suggests a general lack of awareness among participants, both expert and non-expert, about the difficulty of finding the optimal solution. As suggested by the literature, experts are not immune to biases (Berthet, 2022; Blumenthal-Barby & Krieger, 2015; Hodgkinson et al., 1999; Kumar & Goyal, 2015; Lidén et al., 2019).

The exploration bias towards favourites seems to stem from a cognitive factor independent of confirmation bias. The correlations between confirmation bias tasks and exploration bias are close to zero, meaning the empirical result is clear. We can conclude that the biases in the task are not driven by confirmation bias. We realise that the confirmation bias tasks and the set selection task involve different cognitive processes. As stated in the introduction, confirmation bias tasks require participants to seek information to verify a logical implication. This commonality explains why previous studies have shown that these tasks align with a common cognitive factor (Berthet, Teovanovic, et al., 2022; Klayman & Ha, 1987). However, the set selection paradigm is constructed differently, without a logical implication to verify. Exploration in the set selection task requires seeking the most probabilistically useful information to then make a multi-selection decision. The two-step reasoning, the computational dimension, and the need to make decisions based on sets are important parameters of the set selection task, distinguishing it clearly from the confirmation bias tasks tested in Exp. 4.

The Mindware Gaps factor identified in the literature was found in our analyses and helps explain the mathematical errors made during the selection process. Participants need to apply probabilistic reasoning to avoid selection errors, as revealed by exploratory factor analyses. The BNT and CRT tasks were included in the Mindware Gaps factor of previous taxonomies, so it makes sense that they appear in the same factor (Ceschi et al., 2019; Stanovich et al., 2008). The CRT and BNT tasks, designed to assess abilities rather than biases, logically contribute to Mindware Gaps. Numeracy-related abilities and biases thus influence overall selection errors in the set selection paradigm.

### c. Limitations of the Study

Our study highlights the challenges of replicating robust cognitive bias factors across different tasks. The lack of a clear link between bias tasks and our paradigm suggests that other taxonomies or bias tasks may be necessary to uncover consistent factors that explain exploration and selection biases. As taxonomies of cognitive factors are recent, it is still a challenge to find an established consensus on which to build an experiment.

A larger sample size would allow for more reliable factor analyses, providing stronger insights into the cognitive underpinnings of performance in this task. Ceschi et al. (2019) and Berthet et al. (2022) performed factorial analyses with 434 and 223 participants, respectively, while we had 83 participants. A larger sample would also allow us to exclude those who seem to perform the task randomly. As we used an original paradigm, it was difficult to define criteria to identify participants who were not taking the task seriously. For instance, exploring randomly is a perfectly acceptable strategy within the set selection paradigm. In future research, the combination of data on trial completion time and the number of mathematical errors should enable us to filter such participants more effectively.

Regarding recruitment expertise, it is possible that the recruiters we tested are not the ideal experts for this paradigm. Although recruiters are accustomed to creating shortlists based on gathered information, their task typically culminates in the selection of a single candidate. They do not necessarily need to make explicit choices between information before deciding.

### d. Future Directions and Conclusion

A more suitable expert group might be programme directors or other professionals who must manage large pools of candidates to form cohorts, as their experience aligns more closely with the two-step multi-selection reasoning required in our task. It may also be worthwhile

to focus on recruitment experts involved in sourcing or headhunting, who are used to being at the early stages of the recruitment process.

To make the task more realistic for professionals, we could adapt it so that exploration is implicitly constrained by financial or time costs, rather than limiting actions. As noted by (Berthet, 2022), even if experts display biases in theoretical tasks, it is unclear whether these biases manifest in their professional work. Making the task more realistic would help bridge the gap between lab studies and real-world experience.

Another interesting result to investigate is the confidence reported by participants in their exploration strategy. While the confidence reported after choosing a global strategy is generally moderate and does not predict its quality, what about trial-by-trial confidence? Measuring confidence after each trial would allow us to gain more insight into participants' perceptions of how effectively they apply their strategy, trial by trial. This would allow us to explore their metacognition, i.e., their knowledge about their own knowledge (Ackerman & Thompson, 2017).

In conclusion, while our study has provided valuable insights into the role of cognitive biases and expertise in decision-making, it also highlights the complexity of the set selection task. Our findings encourage further exploration of this paradigm to better understand why participants struggle with the task and why they approach it in such divergent ways. One promising avenue for future research is to develop interventions that help participants improve their performance. By testing various remediation strategies, we can identify the cognitive mechanisms that disturb performance and determine whether targeted support can help participants overcome these challenges.

## Bibliography

1. Abbas, S. I., Shah, M. H., & Othman, Y. H. (2021). Critical Review of Recruitment and Selection Methods: Understanding the Current Practices. *Annals of Contemporary Developments in Management & HR (ACDMHR)*, 3(3), Article 3. <https://doi.org/10.33166/ACDMHR.2021.03.005>
2. Ackerman, R., & Thompson, V. A. (2017). Meta-Reasoning: Monitoring and Control of Thinking and Reasoning. *Trends in Cognitive Sciences*, 21(8), 607–617. <https://doi.org/10.1016/j.tics.2017.05.004>
3. Aczel, B., Bago, B., Szollosi, A., Foldes, A., & Lukacs, B. (2015). Measuring Individual Differences in Decision Biases: Methodological Considerations. *Frontiers in Psychology*, 6. <https://doi.org/10.3389/fpsyg.2015.01770>
4. Arnott, D. (2006). Cognitive biases and decision support systems development: A design science approach. *Information Systems Journal*, 16(1), 55–78. <https://doi.org/10.1111/j.1365-2575.2006.00208.x>
5. Berthet, V. (2021). The Measurement of Individual Differences in Cognitive Biases: A Review and Improvement. *Frontiers in Psychology*, 12. <https://doi.org/10.3389/fpsyg.2021.630177>
6. Berthet, V. (2022). The Impact of Cognitive Biases on Professionals' Decision-Making: A Review of Four Occupational Areas. *Frontiers in Psychology*, 12. <https://doi.org/10.3389/fpsyg.2021.802439>
7. Berthet, V., Autissier, D., & De Gardelle, V. (2022). Individual differences in decision-making: A test of a one-factor model of rationality. *Personality and Individual Differences*, 189, 111485. <https://doi.org/10.1016/j.paid.2021.111485>
8. Berthet, V., Autissier, D., & de Gardelle, V. (2022). Individual differences in decision-making: A test of a one-factor model of rationality. *Personality and Individual Differences*, 189, 111485. <https://doi.org/10.1016/j.paid.2021.111485>
9. Berthet, V., & de Gardelle, V. (2023). The heuristics-and-biases inventory: An open-source tool to explore individual differences in rationality. *Frontiers in Psychology*, 14. <https://www.frontiersin.org/articles/10.3389/fpsyg.2023.1145246>
10. Berthet, V., Teovanovic, P., & de Gardelle, V. (2022). *Confirmation bias in hypothesis testing: A unitary phenomenon?* [Preprint]. Open Science Framework. <https://doi.org/10.31219/osf.io/wjkr5>
11. Blumenthal-Barby, J. S., & Krieger, H. (2015). Cognitive Biases and Heuristics in Medical Decision Making: A Critical Review Using a Systematic Search Strategy. *Medical Decision Making*, 35(4), 539–557. <https://doi.org/10.1177/0272989X14547740>
12. Bruine De Bruin, W., Parker, A. M., & Fischhoff, B. (2007). Individual differences in adult decision-making competence. *Journal of Personality and Social Psychology*, 92(5), 938–956. <https://doi.org/10.1037/0022-3514.92.5.938>
13. Burgoyne, A. P., Mashburn, C. A., Tsukahara, J. S., Hambrick, D. Z., & Engle, R. W. (2023). Understanding the relationship between rationality and intelligence: A latent-variable approach. *Thinking & Reasoning*, 29(1), 1–42. <https://doi.org/10.1080/13546783.2021.2008003>
14. Ceschi, A., Costantini, A., Sartori, R., Weller, J., & Di Fabio, A. (2019). Dimensions of decision-making: An evidence-based classification of heuristics and biases. *Personality and Individual Differences*, 146, 188–200. <https://doi.org/10.1016/j.paid.2018.07.033>
15. Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97. <https://doi.org/10.1016/j.jbef.2015.12.001>
16. Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring Risk Literacy: The Berlin Numeracy Test. *Judgment and Decision Making*, 7(1), 25–47. <https://doi.org/10.1017/S1930297500001819>

17. Cole, M. S., Rubin, R. S., Feild, H. S., & Giles, W. F. (2007). Recruiters' Perceptions and Use of Applicant Résumé Information: Screening the Recent Graduate. *Applied Psychology*, 56(2), 319–343. <https://doi.org/10.1111/j.1464-0597.2007.00288.x>
18. De Neys, W., & Glumicic, T. (2008). Conflict monitoring in dual process theories of thinking. *Cognition*, 106(3), 1248–1299. <https://doi.org/10.1016/j.cognition.2007.06.002>
19. Dimara, E., Franconeri, S., Plaisant, C., Bezerianos, A., & Dragicevic, P. (2020). A Task-Based Taxonomy of Cognitive Biases for Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 26(2), 1413–1432. [IEEE Transactions on Visualization and Computer Graphics. https://doi.org/10.1109/TVCG.2018.2872577](https://doi.org/10.1109/TVCG.2018.2872577)
20. D'Silva, C. (2020). A Study On Increase in E-Recruitment and Selection Process. *International Journal of Research in Engineering, Science and Management*, 3(8), Article 8.
21. Erceg, N., Galić, Z., & Bubić, A. (2022). Normative responding on cognitive bias tasks: Some evidence for a weak rationality factor that is mostly explained by numeracy and actively open-minded thinking. *Intelligence*, 90, 101619. <https://doi.org/10.1016/j.intell.2021.101619>
22. Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, 19(4), 25–42. <https://doi.org/10.1257/089533005775196732>
23. Hodgkinson, G. P., Bown, N. J., Maule, A. J., Glaister, K. W., & Pearman, A. D. (1999). Breaking the frame: An analysis of strategic cognition and decision making under uncertainty. *Strategic Management Journal*, 20(10), 977–985. [https://doi.org/10.1002/\(SICI\)1097-0266\(199910\)20:10<977::AID-SMJ58>3.0.CO;2-X](https://doi.org/10.1002/(SICI)1097-0266(199910)20:10<977::AID-SMJ58>3.0.CO;2-X)
24. Holt, C. A., & Laury, S. K. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 92(5), 1644–1655. <https://doi.org/10.1257/000282802762024700>
25. Jacowitz, K. E., & Kahneman, D. (1995). Measures of Anchoring in Estimation Tasks. *Personality and Social Psychology Bulletin*, 21(11), 1161–1166. <https://doi.org/10.1177/01461672952111004>
26. Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237.
27. Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, 39(4), 341–350. <https://doi.org/10.1037/0003-066X.39.4.341>
28. Kassambara, A. (2023). *Ggpubr: 'ggplot2' Based Publication Ready Plots. R package version 0.6.0.* <https://CRAN.R-project.org/package=ggpubr>
29. Klayman, J. (1995). Varieties of Confirmation Bias. In J. Busemeyer, R. Hastie, & D. L. Medin (Eds.), *Psychology of Learning and Motivation* (Vol. 32, pp. 385–418). Academic Press. [https://doi.org/10.1016/S0079-7421\(08\)60315-1](https://doi.org/10.1016/S0079-7421(08)60315-1)
30. Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94(2), 211–228. <https://doi.org/10.1037/0033-295X.94.2.211>
31. Kumar, S., & Goyal, N. (2015). Behavioural biases in investment decision making – a systematic literature review. *Qualitative Research in Financial Markets*, 7(1), 88–108. <https://doi.org/10.1108/QRFM-07-2014-0022>
32. Lidén, M., Gräns, M., & Juslin, P. (2019). 'Guilty, no doubt': Detention provoking confirmation bias in judges' guilt assessments and debiasing techniques. *Psychology, Crime & Law*. <https://www.tandfonline.com/doi/full/10.1080/1068316X.2018.1511790>
33. Long, J. A. (2024). jtools: Analysis and Presentation of Social ScientificData. *Journal of Open Source Software*, 9(101), 6610. <https://doi.org/10.21105/joss.06610>
34. Lüdecke, D., S. Ben-Shachar, M., Patil, I., M. Wiernik, B., Bacher, E., Thériault, R., & Makowski, D. (2022). *easystats: Framework for Easy Statistical Modeling, Visualization, and Reporting.* CRAN. <https://easystats.github.io/easystats/>

## Chapter 2: Cognitive Factors Influencing Information Search Biases

35. O. Wilke, C. (2024). *Cowplot: Streamlined Plot Theme and Plot Annotations for 'ggplot2'*. R package version 1.1.3. <https://CRAN.R-project.org/package=cowplot>
36. R Core Team. (2023). *R: A Language and Environment for Statistical Computing*. <https://www.R-project.org/>
37. Roulin, N., & Levashina, J. (2019). LinkedIn as a new selection method: Psychometric properties and assessment approach. *Personnel Psychology*, 72(2), 187–211. <https://doi.org/10.1111/peps.12296>
38. Snyder, M., & Swann, W. B. (1978). Hypothesis-testing processes in social interaction. *Journal of Personality and Social Psychology*, 36, 1202–1212. <https://doi.org/10.1037/0022-3514.36.11.1202>
39. Stanovich, K. E., Toplak, M. E., & West, R. F. (2008). The Development of Rational Thought: A Taxonomy of Heuristics and Biases. In R. V. Kail (Ed.), *Advances in Child Development and Behavior* (Vol. 36, pp. 251–285). JAI. [https://doi.org/10.1016/S0065-2407\(08\)00006-2](https://doi.org/10.1016/S0065-2407(08)00006-2)
40. Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology*, 94(4), 672–695. <https://doi.org/10.1037/0022-3514.94.4.672>
41. Teovanović, P., Knežević, G., & Stankov, L. (2015). Individual differences in cognitive biases: Evidence against one-factor theory of rationality. *Intelligence*, 50, 75–86. <https://doi.org/10.1016/j.intell.2015.02.008>
42. Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315. <https://doi.org/10.1037/0033-295X.90.4.293>
43. Wason, P. C. (1960). On the Failure to Eliminate Hypotheses in a Conceptual Task. *Quarterly Journal of Experimental Psychology*, 12(3), 129–140. <https://doi.org/10.1080/17470216008416717>
44. Wason, P. C. (1968). Reasoning about a Rule. *Quarterly Journal of Experimental Psychology*, 20(3), 273–281. <https://doi.org/10.1080/14640746808400161>
45. Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Golemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., ... Yutani, H. (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, 4(43), 1686. <https://doi.org/10.21105/joss.01686>
46. Zeileis, A., Meyer, D., & Hornik, K. (2007). Residual-Based Shadings for Visualizing (Conditional) Independence. *Journal of Computational and Graphical Statistics*, 16(3), 507–525. <https://doi.org/10.1198/106186007X237856>

## Supplementary materials

### The Questionnaire Version Replicates the Original Task Version

#### Methods

In Exp. 4, after completing 15 trials of the original version of the main task, participants completed a questionnaire version of the task (MTq). MTq is a theoretical version of the same task, without any financial incentives for performance. They answered a series of questions to assess their exploration strategy during the main task. First, they indicated whether they believed an optimal exploration strategy existed (using a 5-point Likert scale). The following question is the focus of this section: participants had to choose what they considered to be the most optimal exploration strategy from among Highest, Lowest, Middle, Extreme, Random, or a different strategy for each trial. Response options were presented in a random order. After making their choice, participants indicated their level of confidence in their answer (from 1 = I guessed to 4 = I am certain) and the frequency with which they used the indicated strategy (1 = For none of the trials to 4 = For all the trials). Next, they had to provide a secondary response, which was the exploration strategy they believed would be optimal if the one they initially indicated was not actually the most optimal. They also indicated their confidence and the frequency with which they used this secondary strategy.

In Exp. 5, participants had a presentation of the problem to understand it. Contrary to Exp. 4 participants', they did not perform the original task before to answer the questionnaire. The task presentation is represented on Fig. 7.

### Description du problème

Vous avez un ensemble de **10 profils de candidat-es** à consulter dans le cadre d'une campagne de recrutement pour choisir **les 5 meilleur-es candidat-es**. Chaque candidat-e a réalisé **2 évaluations de compétences** :

- Une évaluation d'un ensemble A de compétences, représenté par **un score A**.
- Une évaluation d'un ensemble B de compétences (distinct de l'ensemble A) représenté par **un score B**.

Les scores A sont compris **entre 0 et 10** et sont distribués de manière **uniforme**.

Les scores B sont compris **entre 0 et 10** et sont distribués de manière **uniforme**.

Le score A et le score B d'un-e candidat-e sont **indépendants**.

Plus un score est **élevé**, plus le-la candidat-e est évalué-e comme étant **compétent**.

Au départ, seul le score A de chaque candidat-e est révélé (c'est-à-dire visible pour vous).

Le recrutement se déroule en **2 étapes** :

1. Explorer le score B (c'est-à-dire voir le score B) de **5 candidat-es** de votre choix parmi les 10 candidat-es.
2. Choisir les **5 meilleur-es candidat-es parmi les 10 candidat-es**, à savoir les 5 candidat-es ayant la meilleure moyenne en considérant leurs scores A et B.

**Attention** : Les scores A et B de chaque candidat-e sont pris en considération pour déterminer les meilleur-es candidat-es, que vous ayez choisi d'explorer leur score B ou non !

Nous allons vous poser deux questions sur le problème qui vient de vous être énoncé. Appuyer sur "Suivant" lorsque vous êtes prêt-e.

Fig. 7. Screenshot of Task presentation in Exp. 5.

## Results

### *Within replication of the task original version with the questionnaire*

In Experiment 4, in the short version, 57% of participants believe the optimal exploration strategy is *Highest* (vs. 73% in the original version), with *Middle* being the second most chosen strategy (22% vs. 22% in the original version), which is the most optimal heuristic. We observe a slight difference in distribution between the practical versions of the task and the questionnaire version ( $\chi^2 = 18.26$ ,  $p = .019$ ; Adjusted Cramer's  $V = .15$ ). Specifically, within Experiment 4, there is a slight but significant decrease in the proportion of participants choosing *Highest* in the questionnaire compared to those categorised as *Highest* during their 15 trials ( $\chi^2 = 4.48$ ,  $p = .034$ ; Adjusted Cramer's  $V = .16$ ). However, this decrease is not associated with a significant increase in the proportion of participants choosing the more optimal heuristic (*Middle*) compared to those near *Middle* during the 15 trials of the main task ( $\chi^2 = .30$ ,  $p = .586$ ). Indeed, participants are somewhat more evenly distributed across

incorrect responses (*Highest, Lowest, Extreme* primarily) in the questionnaire version compared to the practical version.

When comparing responses between the two versions, 77% of participants in the questionnaire chose the heuristic they were closest to in the practical version. All participants reported that they effectively used the heuristic they chose in the questionnaire version in their 15 trials (with 98% of participants reporting using this strategy for more than half of the trials). We find that the overlap between the most applied heuristic is higher when it is also the one chosen in the questionnaire, compared to when the most applied heuristic is not chosen in the questionnaire version ( $W = 308.50$ ,  $p < .001$ ;  $r = -.49$ , 95% CI =  $[-.68, -.24]$ ). Additionally, reaction time does not have a significant link with the quality of the response chosen in the questionnaire ( $W = 532.00$ ,  $p = .914$ ). Confidence in their response (expressed on a scale from 1: I don't know to 4: I am certain) also does not have a significant link with the quality of the response ( $W = 718.50$ ,  $p = .222$ ). Confidence is not correlated with response time to choose a strategy ( $r = -.09$ ,  $p = .416$ ) either.

#### *Between-Experiments Replication*

This questionnaire version of the task produces a stable proportion of optimal participants compared to the results of Experiment 4. The proportion of participants whose first choice is the best heuristic is 24% in Experiment 5 and 27% in Experiment 4 ( $\chi^2 = .09$ ,  $p = .764$ ). Similarly, across the 15 trials of the task in Experiment 4, the proportion of optimal participants (those using the best heuristic) was 22%, similar to that observed in Experiment 5 ( $\chi^2 = .03$ ,  $p = .868$ ). In Experiment 1, the proportion was also similar (20%,  $\chi^2 = .15$ ,  $p = .700$ ).

The most chosen strategy remains to explore the favourites, although not always in the same proportions (Fig. 8). In the Experiment 5 questionnaire, 38% of participants chose the Highest heuristic (as their first choice), compared to 57% in the Experiment 4 questionnaire ( $\chi^2 = 6.33$ ,  $p = .012$ ; Adjusted Cramer's V = .17); 73% in the original task version of Experiment 4 ( $\chi^2 = 24.29$ ,  $p < .001$ ; Adjusted Cramer's V = .35); and 65% in the original task version of Experiment 5 ( $\chi^2 = 13.99$ ,  $p < .001$ ; Adjusted Cramer's V = .27).

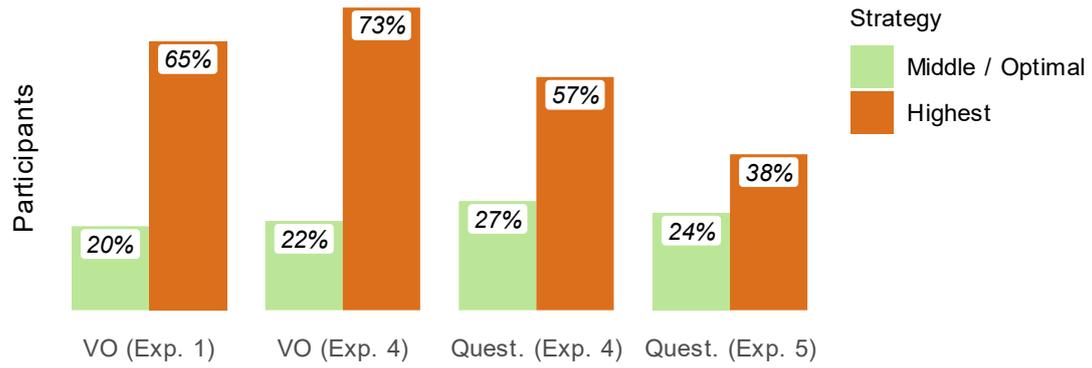


Fig. 8. Graph representing the proportion of participants using the Highest and Middle strategies across experiments and task formats. VO = original version of the task. The VO contained 25 trials in Experiment 1 and 15 trials in Experiment 4. In the VO, use of the optimal strategy is grouped with the Middle heuristic. Quest. = questionnaire version with a single-choice question. In this version, the most optimal strategy offered was Middle.

# Chapter 3: Metacognitive Judgements in Exploration Strategies for Set Selection

Jean-Michel Dagba<sup>1,2</sup>, Alexandre Lietard<sup>3,4</sup>, Jean-Christophe Vergnaud<sup>1,5</sup>, Vincent de Gardelle<sup>1,5,6</sup>

<sup>1</sup> Centre d'Économie de la Sorbonne (CES) & Université Paris 1 Panthéon Sorbonne, France

<sup>2</sup> Entreprise Humans Matter, France

<sup>3</sup> Université Paris Cité, France

<sup>4</sup> KU Leuven, Belgium

<sup>5</sup> Centre national de la recherche scientifique (CNRS), France

<sup>6</sup> Paris School of Economics (PSE), France

## Table of contents

Table of contents	97
I. Introduction	98
II. Method	100
a. Procedure	100
b. Measures	101
c. Participants	103
d. Apparatus	103
e. Statistical analysis	103
III. Results	103
a. Block 1 Replicates Previous Experiments	103
b. Similar pattern and biases between block 1 and 2, but with slight improvements	104
c. Participants overestimate the optimality of their exploration	105
d. Exploration confidence is based on exploration self-consistency	106
e. Faster participants are more confident	107
f. Overestimation is not influenced by socio-demographic characteristics	108
IV. Discussion	109
a. Study summary and main findings	109
b. Exploration confidence is driven by self-consistency, not by optimality	109
c. Future Directions	110
Bibliography	112

## I. Introduction

In Chapter 2, we found that participants' confidence in their exploration strategy choice in the questionnaire version was not predictive of its quality. People who choose a wrong strategy were as confident as people choosing the good strategy. This highlights a gap in how well individuals perceive the quality of their exploration strategy in the questionnaire version, underscoring the importance of studying confidence of the exploration strategy in the original version of the task. This initial result raises further questions about participants' ability to recognise whether their exploration strategy is effective.

In this chapter, we measure trial-by-trial confidence in the exploration strategy within the original version of the set selection paradigm to understand participants' metacognition in such a complex reasoning task. Metacognition, as defined by Nelson & Narens (1990), refers to the ability to monitor and evaluate one's cognitive processes and adjust efforts accordingly. Metareasoning, a subset of metacognition, focuses specifically on the monitoring and regulation of reasoning and problem-solving activities (Ackerman & Thompson, 2017). On the one hand, measuring confidence allows us to replicate the confidence results from Chapter 2, comparing participants' confidence with the actual optimality of their strategy. On the other hand, trial-by-trial measures offer novel insights into participants' metacognitive sensitivity, or their ability to accurately discriminate between correct and incorrect responses (Fleming & Lau, 2014). While metacognitive sensitivity has been extensively studied in domains such as memory and perception (Jin et al., 2022), there is a lack of research on its application to complex reasoning tasks like the set selection paradigm, making strong predictions difficult.

The focus of this chapter is to explore what participants' confidence can reveal about their exploration strategy. Does their confidence provide meaningful insight into the quality of their exploration? If so, what factors is their confidence calibrated to? Could their metacognitive sensitivity reflect their understanding of the optimal strategy, or is it more closely aligned with their preferred, but possibly suboptimal, strategy? Does their confidence allow them to express their competence more accurately than their exploration strategy itself?

To evaluate participants' confidence, they were informed that it is possible to optimise exploration to maximise their chances of success in the selection stage. This information gave them the opportunity to use their confidence judgements to express insights beyond their direct exploration behaviour.

It is likely that participants overestimate the quality of their exploration strategy in this task. Participants tend to deliberately explore the favourites without recognising that this approach is biased. Overconfidence is a well-documented cognitive bias category that includes three phenomena: overestimation, i.e., overestimating one's performance; overplacement, i.e., believing oneself to be better than others; and overprecision, i.e., being overly confident in the accuracy of one's judgements (Duttle, 2016; Griffin & Tversky, 1992; Lichtenstein et al., 1982). Individuals tend to show overconfidence in difficult tasks but may display underconfidence in easier tasks (Moore & Healy, 2008). These phenomena suggest that individuals tend to poorly assess the overall situation and evaluate the difficulty of a task. The experiment in this chapter allows us to determine whether participants overestimate or underestimate the optimality of their exploration. This may reflect the perceived difficulty of the task.

By measuring metacognitive sensitivity, we can determine whether participants are more attuned to their preferred strategy or to the actual optimality of their strategy. For instance, this could help reveal whether a participant close to the *Highest* heuristic adjusts their confidence based on their similarity to the *Highest* strategy or based on how closely they align with the *Optimal* exploration strategy. We aim to detect the direction of their metacognitive sensitivity.

As previously mentioned, participants do not seem to identify the bias in their strategy and, therefore, their distance from the optimal one. It is likely that their confidence is calibrated to the strategy they consistently apply, rather than the optimal strategy. Following the principle of self-consistency (Koriat, 2012), the more consistently they use a particular exploration strategy, the more confident they may become. This self-consistency may manifest both across participants—where those closer to their preferred strategy exhibit higher confidence—and within individual trials, where a participant's metacognitive sensitivity is calibrated to their preferred strategy, leading to greater confidence in trials that align with that strategy.

Metacognitive sensitivity aligned with participants' preferred strategy, combined with overconfidence in the quality of their exploration, may reduce their motivation to engage more cognitive resources to find a better strategy. If participants possess good metacognitive calibration, their sensitivity will be driven by how closely their strategy aligns with the optimal one. In such cases, they could estimate that the quality of their exploration is suboptimal and adjust by seeking a better strategy.

## II. Method

### a. Procedure

Participants engage in an experiment lasting approximately 40 minutes. The procedure, as shown in Fig. 1, consists of three main stages: (1) A block of 15 trials of the set exploration and selection task, as in Exp. 4, (2) A new block of 15 trials of the same task, with confidence judgements following each exploration, and (3) A socio-demographic questionnaire.

We retained a first control block of 15 trials to ensure that participants' performance was similar to that of previous studies and to verify that their behaviour in block 2 was not influenced by the introduction of quality estimations of their explorations.

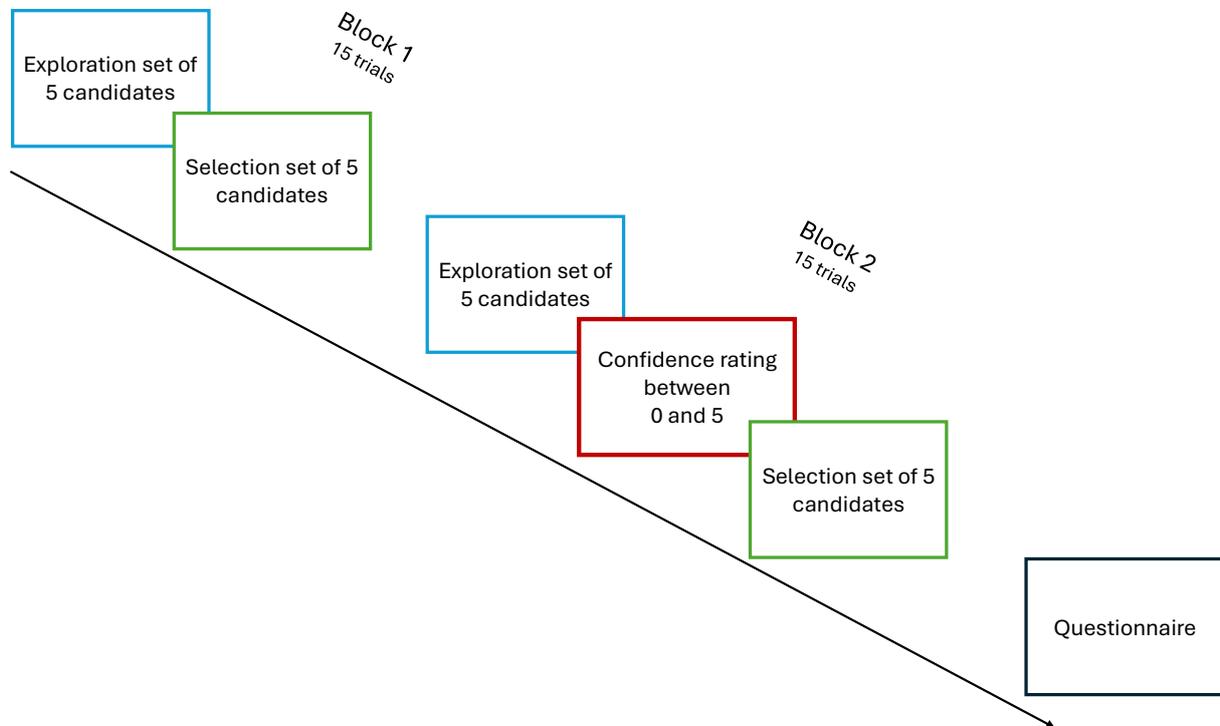


Fig. 1.

Schéma des étapes de la procédure de l'expérience 6.

### Main task

During the first block of trials, participants perform the standard version of the task (see details in Exp. 1 of Chapter 1 or Exp. 4 of Chapter 2), with no differences in task instructions or any particular warning about the modifications to be introduced in the following block. The only difference is that participants complete 15 trials (instead of 25 trials).

## **Main task with estimation reports**

The second block of trials involves the same task, with the addition of an estimation of the quality of their exploration, to be reported immediately after the exploration on a scale from 0 to 5. This part begins with instructions on how to make an estimation after the exploration phase (and before the selection phase). Participants must evaluate how many of the candidates they chose to explore were among the most useful ones to explore. The instructions explain that it is possible to determine to what extent a candidate is useful to explore (in order to find the top 5 candidates in the selection phase). Thus, candidates can be ranked from most useful to least useful to explore. Some candidates may be tied, meaning there will always be at least 5 candidates more useful to explore than the others (see Supplementary Materials for screenshots of the instructions). The minimal differences between blocks 1 and 2 indicate that the introduction of this additional step did not alter the way participants performed the task, allowing us to gather reliable data on their metacognition. As in the previous block, participants have a training trial before completing the 15 trials.

### **b. Measures**

We use the same measures as those used in Chapters 1 and 2, except for those related to confidence, which are detailed below. For measures of optimal exploration, optimal selection, and performance, please refer to the methods in Chapter 1.

#### **Optimal exploration**

Optimal exploration refers to the set of candidates to explore in order to maximise expected gains in a given trial (see details in Chapters 1 or 2).

#### **Preferred strategy**

The preferred strategy of a participant is the strategy they are closest to, on average, among the strategies: Highest (exploring the 5 candidates with the highest A scores), Lowest (the lowest A scores), Extreme (the very highest and lowest scores), Random (random exploration), Middle (neither the highest nor lowest), and the Optimal strategy (as defined in Chapter 1).

#### **Exploration optimality**

Exploration optimality is the number of candidates a participant actually explored in alignment with the optimal exploration strategy. If two candidates are tied (i.e. if they have

the same A scores), exploring either is considered equally optimal. This is calculated for every trial (to allow within-participant analysis), averaged, and expressed as a percentage to obtain a mean value for each participant (allowing across-participant analysis).

### **Exploration consistency**

Exploration consistency is the number of candidates a participant explored in line with their preferred strategy. If two candidates are tied, exploring either is considered equivalent. This is calculated for every trial and averaged.

### **Exploration confidence**

The confidence is the number of candidates that a participant believes they have explored according to the optimal exploration strategy. It is recorded for each trial to provide within-participant measures. This is also averaged and expressed as a percentage to obtain a mean value for each participant and allow across-participant comparison.

### **Overconfidence (overestimation)**

Overestimation is the difference between confidence in exploration optimality and actual exploration optimality. It is expressed in percentage points. A positive value corresponds to overestimation of optimality, while a negative value indicates underestimation.

We also computed the difference between confidence in exploration and exploration consistency. A positive value corresponds to overestimation of exploration consistency, while a negative value indicates underestimation.

### **Metacognitive sensitivity**

Metacognitive sensitivity is measured within participants by calculating the correlation between confidence and overlap with the optimal strategy.

### **Exploration time**

Exploration time is measured for each trial in milliseconds, starting from the appearance of the screen until the last click on a candidate.

### **Confidence response time**

The time taken to report confidence in exploration is measured for each trial in milliseconds, starting from the appearance of the screen until the participant submits their confidence response.

## Optimal selection

Optimal selection refers to the set of candidates to be selected to maximise expected gains, given a defined exploration (see details in Chapter 1 or 2).

## Performance

Performance is the number of selected candidates that are indeed among the best. This is averaged across all trials and expressed as a percentage (see details in Chapter 1 or 2).

### c. Participants

Experiment 6 involved 78 participants (41 women). Participants' mean age was 26 years (ranging from 18 to 58). The mean gain was 9€ added to a participation fee of 4€.

### d. Apparatus

Experiments were conducted online. They were build using oTree (Chen et al., 2016) with the technical support of the Fédération S2CH.

### e. Statistical analysis

Statistical analyses were conducted using R (R Core Team, 2023), and the following packages: 'tidyverse' (Wickham et al., 2019), 'ggpubr' (Kassambara, 2023), 'cowplot' (O. Wilke, 2024) and 'easystats' (Lüdecke et al., 2022).

## III. Results

### a. Block 1 Replicates Previous Experiments

All results from previous experiments were replicated in the first Control block of 15 trials in this experiment. The average performance was 83%, which was suboptimal ( $W = 2901.00$ ,  $p < .001$ ;  $r = .98$ ). The average performance loss due to exploration was 5.47 percentage points ( $W = 2691.00$ ,  $p < .001$ ;  $r = .99$ ), and the loss due to selection was 3.42 percentage points ( $W = 2273.50$ ,  $p < .001$ ;  $r = .78$ ). For 67% of participants the preferred strategy was the *Highest* heuristic, and only 21% were optimal or used the *Middle* heuristic. On average, participants' exploration was 57% aligned with optimal exploration. Consequently, there was still a significant exploration bias towards favourites ( $W = 2495.50$ ,  $p < .001$ ;  $r = .71$ ). Similarly, participants continued to show a selection bias towards explored options ( $W = 2168.50$ ,  $p <$

.001;  $r = .56$ ) and a significant performance loss of about 1% due to mathematical errors ( $W = 1128.00$ ,  $p < .001$ ;  $r = .64$ ).

### b. Similar pattern and biases between block 1 and 2, but with slight improvements

Participants performed the Block 2 (with estimations on exploration) consistently with their Block 1 performances, as shown in Figure 2. The average exploration bias towards favourites remained stable between blocks ( $W = 1238.50$ ,  $p = 0.279$ ), as did the selection bias towards explored options ( $W = 1005.00$ ,  $p = .118$ ). However, the average exploration optimality improved from Block 1 (mean = +2.5 percentage points:  $W = 468.50$ ,  $p = .004$ ;  $r = .43$ ), as did the average selection optimality (mean = +1.78 percentage points:  $W = 530.50$ ,  $p = .002$ ;  $r = .46$ ) and the final performance (mean = +2.02 percentage points:  $W = 463.50$ ,  $p < .001$ ;  $r = .54$ ). Analyses indicated a consistency of 76% in the strategy within individuals, largely exceeding chance level of 15%).<sup>1</sup> Overall, participants' exploration strategy was not altered by the introduction of instructions regarding the self-assessment of exploration or the revelation of the existence of an optimal strategy. Therefore, the confidence estimates can be analysed and interpreted as reliable for task performance in general.

---

<sup>1</sup> As we wanted to check if people are less biased toward the favourites and more optimal, we grouped some strategies to simplify this analysis. We performed the comparison between blocks with the categories Highest vs. Optimal (Optimal or Middle) vs. Other (Lowest, Extreme or Random). With all strategies (Optimal, Middle, Highest, Lowest, Extreme, Random). With all strategies considered, when the best-matching strategies was identical for 56% of participants between blocks (chance level of 22%).

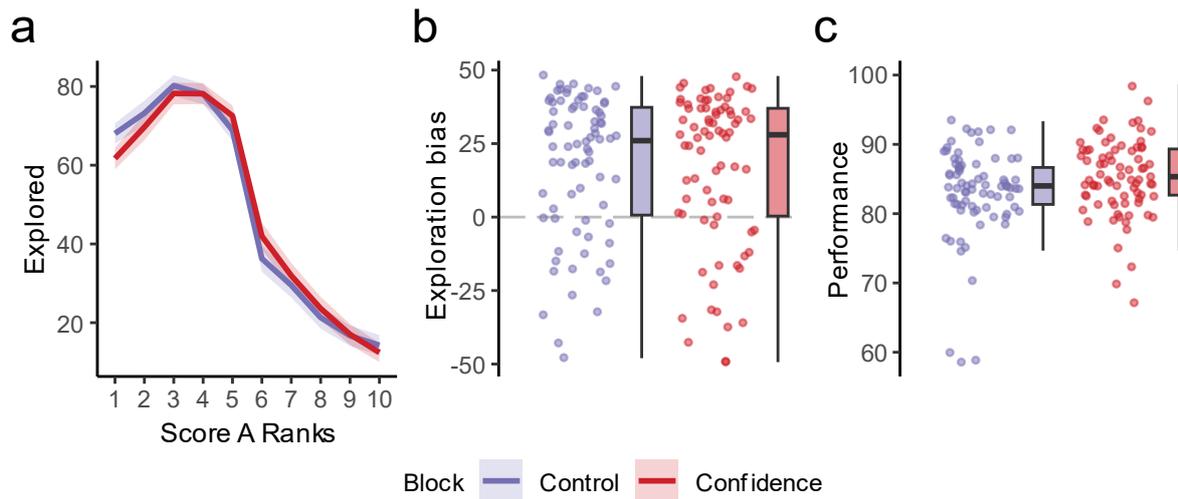


Fig. 2.

Graphs representing the changes between Block 1 (Control) and Block 2 (Exploration Estimation). Panel (a) shows the average exploration strategy of participants in Block 1 (in purple) and Block 2 (in red). The x-axis indicates the rank of the options based on their A score, and the y-axis represents the average percentage of options explored among all options ranked at a given position. Panel (b) presents the exploration bias towards favourites on average in the two blocks. This is the percentage point difference between the proportion of options whose A score ranks in the top 5 explored by participants and the same proportion explored by the optimal exploration strategy. Panel (c) shows the final performance in both blocks. This is the percentage of selected options that were actually in the top 5.

### c. Participants overestimate the optimality of their exploration

On average, participants overestimated the quality of their exploration. Participants rated their exploration as optimal at 76% on average, while it was optimal at 59%, resulting in a 17-percentage points overestimation ( $W = 2668.50$ ,  $p < .001$ ;  $r = .73$ , 95% CI = [.59, .83]).

Confidence does not vary according to the strategy used by participants (Kruskal-Wallis  $\chi^2 = 4.41$ ,  $p = .354$ ), whereas not all heuristics are equally optimal.

Consequently, the degree of overestimation varied significantly depending on the preferred strategy (Kruskal-Wallis  $\chi^2 = 28.68$ ,  $p < .001$ ;  $r = .37$ ), as shown in Figure 3.

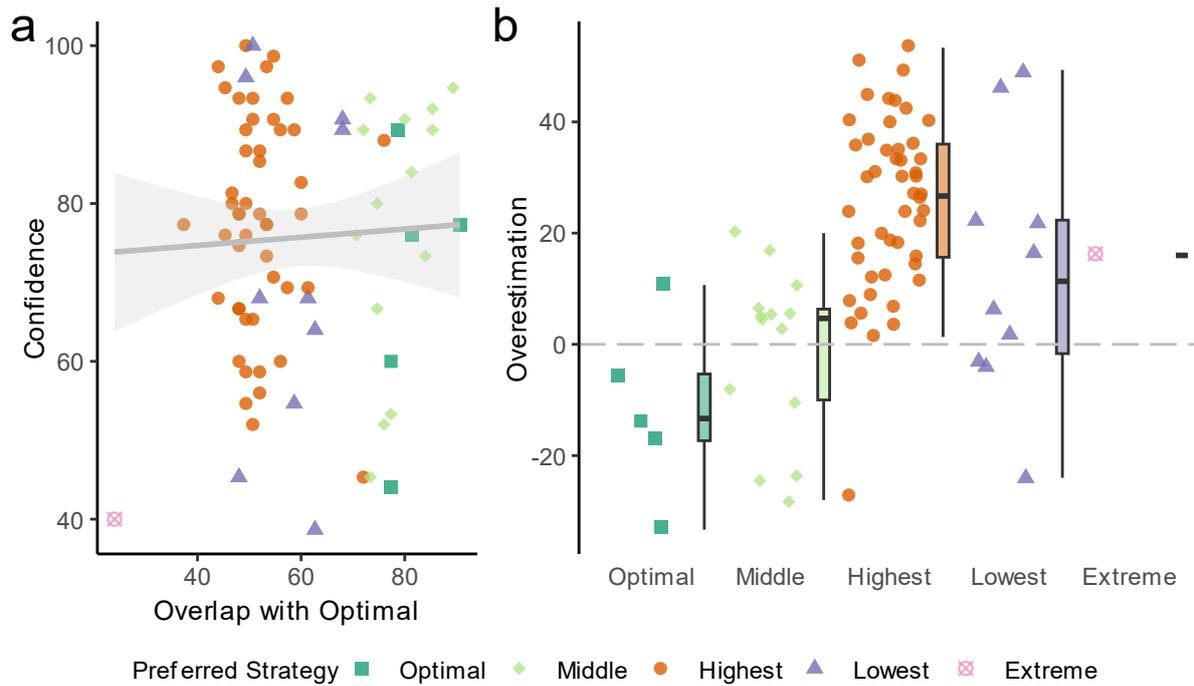


Fig. 3.

Panel (a) illustrates each participant's average estimation of the quality of their exploration, according to the average percentage of exploration optimality. Panel (b) shows the overestimation (in percentage points) based on the preferred strategy of each participant. The overestimation of exploration quality is the difference between the average estimation percentage and the average optimality percentage. The preferred exploration strategy is the strategy that most closely corresponds to the exploration conducted by participants on average, among Optimal, Highest, Middle, Extreme, and Lowest (see Methods for details).

#### d. Exploration confidence is based on exploration self-consistency

In a within-participant analysis, participants do not exhibit good metacognitive sensitivity toward the optimal strategy. Their confidence made after each exploration does not predict the optimality of their exploration in the trial (estimate = .0329,  $p = .291$ ). Across-participants shows no correlation between mean exploration optimality and mean confidence ( $r = .04$ ,  $p = .699$ ).

Participants are, in fact, calibrated to their preferred strategy, and their confidence is consistent with this strategy. Within-participant analysis shows that confidence is predictive of the overlap with the preferred strategy trial by trial (estimate = .055,  $p = .017$ ). Trial-by-trial, their confidence is a little bit sensitive to the proximity to the preferred exploration strategy. Similarly, across-participant analysis shows that participants are calibrated by their preferred strategy. The closer participants are to the strategy that best describes them, the higher their confidence on average ( $r = .41$ ,  $p < .001$ ), as shown in Figure 4.

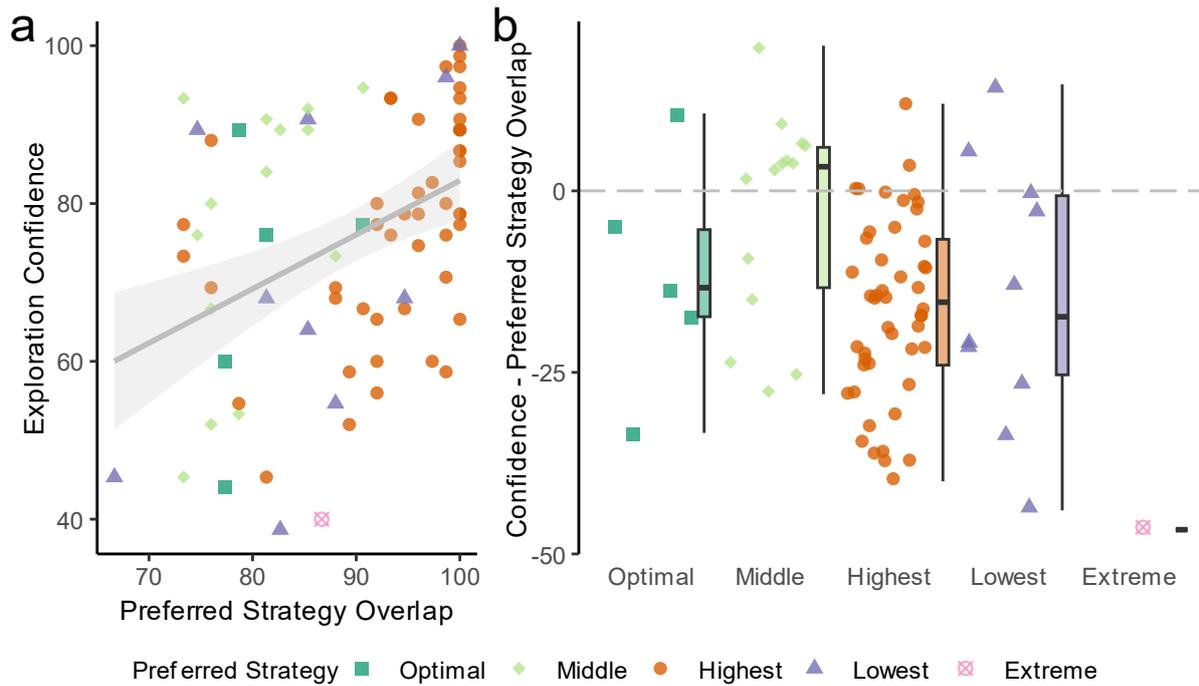


Fig. 4.

Panel (a) illustrates each participant's average estimation of the quality of their exploration, according to the average percentage of overlap with their preferred strategy. Panel (b) shows the difference (in percentage points) between their mean confidence and the mean overlap with their preferred strategy, depending on their preferred strategy.

### e. Faster participants are more confident <sup>2</sup>

The quicker participants are at exploring and estimating, the more confident they are. Across participants, exploration time and the time taken to estimate the quality of one's exploration are strongly correlated ( $r = .45$ , 95% CI = [.24, .61],  $p < .001$ ). Participants who explore quickly also tend to estimate their optimality quickly. The median exploration time is negatively correlated with the average confidence of exploration quality ( $r = -.28$ , 95% CI = [-.48, -.06],  $p = .016$ ) as well as with overestimation ( $r = -.57$ , 95% CI = [-.71, -.39],  $p < .001$ ). Additionally, when the time taken to estimate the quality of exploration increases, the reported estimation significantly decreases ( $r = -.53$ , 95% CI = [-.68, -.34],  $p < .001$ ), and overestimation also decreases ( $r = -.42$ , 95% CI = [-.59, -.21],  $p < .001$ ). In other words, the more time participants take to complete the exploration and self-evaluation, the less they overestimate themselves, as illustrated in Fig. 5.

<sup>2</sup> For across-participants analysis, we excluded 4 participants who were outliers for at least one of the two measures based on time. Outliers are participants whose median time was outside the bounds (1.5 times the interquartile range above the third quartile or below the first quartile) for exploration time or estimation response time.

For within-participants analysis, we excluded 140 trials among the 1170 trials realised in the experiment, with the same exclusion rule applied within-participants response times.

Within participants analysis shows similar results. A participant confidence response time predicts his/her exploration time (estimate = .603,  $p < .001$ ). Faster participant is to complete an exploration stage, more he/she is confident. The more time he/she takes to complete an exploration stage, the less confident he/she is (estimate = -15.140,  $p = .0314$ ). However, taking more time to complete an exploration does not indicate a significant decrease in overconfidence ( $p = .243$ ). Confidence estimation response time is significantly predicted by confidence (estimate = -7.558,  $p < .001$ ) and overconfidence (estimate = -7.393,  $p = .004$ ). The more time a participant takes to assess their confidence following an exploration, the less (over)confident they are for that trial.

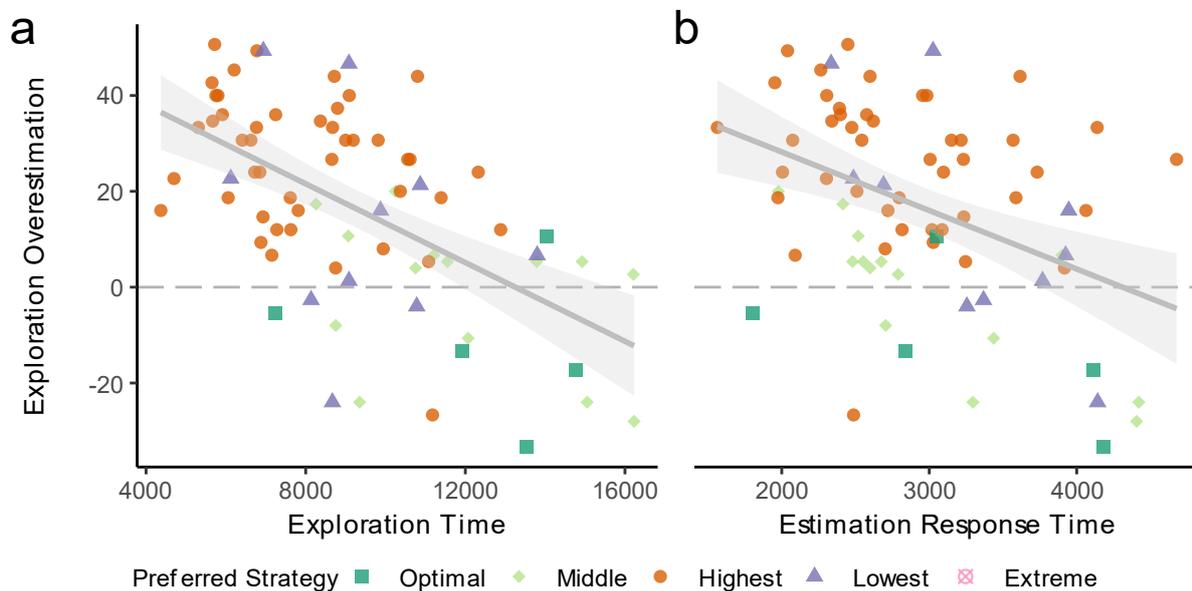


Fig. 5.

Graphs illustrating the correlations between the mean overestimation of exploration quality for each participant depending on (a) the median exploration time and (b) the median confidence response time. The times are in milliseconds. The shapes and colours of the points correspond to the participants' preferred strategy, i.e., the exploration strategy they were closest to on average.

### f. Overestimation is not influenced by socio-demographic characteristics

The overestimation of exploration quality does not appear to be related to socio-demographic characteristics. The analyses show no link with age ( $r = .18$ ,  $p = .107$ ), gender ( $W = 771.00$ ,  $p = .904$ ), level of education (Kruskal-Wallis  $\chi^2 = .71$ ,  $p = .871$ ), or confidence in understanding probabilities ( $r = .14$ ,  $p = .217$ ). There is also no correlation between the number of errors made on the comprehension questions posed in the instructions and the overestimation of exploration quality ( $r = .20$ ,  $p = .077$ ).

## IV. Discussion

### a. Study summary and main findings

This study aimed to explore participants' confidence in their exploration strategy on a trial-by-trial basis within the set selection paradigm. The objective was to investigate the additional information that confidence reveals about participants' exploration strategy.

As in Chapter 2, there is no difference in confidence between optimal and suboptimal participants on average. Consequently, participants who employed an optimal (or near-optimal) exploration strategy made accurate estimates of their performance, whereas those using a suboptimal exploration strategy tended to be overconfident. From an optimality perspective, participants demonstrated poor metacognitive sensitivity on average. The more participants overestimated their exploration optimality, the more they completed the exploration phase quickly and gave rapid confidence estimates. If we consider their preferred strategy instead of the Optimal strategy, participants tend to be more consistent. There was a correlation between participants' confidence and their overlap with the strategy that best characterised them. Participants who applied their preferred strategy more consistently were also the most confident. Trial by trial, a participant was more confident when their exploration aligned more closely with their preferred strategy.

### b. Exploration confidence is driven by self-consistency, not by optimality

All participants express a similar level of confidence, indicating a lack of connection between confidence and optimality. Confidence judgements in reasoning tasks may, therefore, not align with the results presented in studies on metacognition in memory and perception tasks (Jin et al., 2022). Participants with suboptimal exploration strategies seem unaware that their strategy is biased and unsuitable, which explains why they maintain it across trials and blocks. Their sensitivity to the application of a poor heuristic indicates that they likely believe they are using an optimal strategy.

The key difference between optimal (or near-optimal) participants and suboptimal participants lies in the strategy used. While both groups show similar levels of confidence estimates and sensitivity to the proper application of their preferred strategy, only the former group employs the correct strategy.

Participants' consistency in applying their preferred strategy demonstrates their conviction that they have the correct strategy. Their confidence on a trial-by-trial basis also shows a form of metacognitive sensitivity towards applying this preferred strategy. In line with the self-consistency model (Koriat, 2012), the more confident participants are, the more their explorations align with their preferred strategy. Therefore, confidence could be linked to the perception of the consistency of the strategy used rather than the optimality of that strategy.

### c. Future Directions

Previous research has also examined confidence in cognitive bias tasks. Firstly, these studies indicate that individuals report high levels of confidence, similar to our findings – for instance, 72% confidence in Šrol & De Neys (2021). Secondly, these studies align with the theory of conflict detection in bias tasks. They suggest that an individual reports lower confidence when he/she detects a conflict between an intuitive response and a logical-mathematical response in a cognitive bias task compared to when he/she detects no conflict between intuitive and logical-mathematical responses (De Neys, 2010, 2012; De Neys & Bonnefon, 2013; De Neys & Glumicic, 2008; Neys et al., 2011; Šrol & De Neys, 2021). For example, in a conjunction fallacy task, a person named Bill is described as having been good at maths in school. Participants must choose which of two options is more likely to be true. In the conflict version of the task, the response choices are (1) Bill plays in a music band or (2) Bill plays in a music band and is an accountant. The correct answer is option 1 according to probability rules, but the intuitive response is option 2 as the description points to the stereotype that Bill would be an accountant. In the non-conflict version, the description remains the same, but the response options are (1) Bill is an accountant or (2) Bill is an accountant and plays in a music band. Here, option 1 is both the intuitive and the correct answer. In the conflict version, participants report lower confidence in their intuitive (incorrect) response than in their intuitive (correct) response in the non-conflict version. In our study, it can be considered that there is a conflict between the intuitive exploration strategy (Highest) and the correct strategy (Middle). To further investigate individuals' metacognition in this task, one could compare their confidence with a non-conflict version where the intuitive exploration strategy (Highest) is also the optimal strategy (e.g., exploring 5 options to find the top 2). If individuals are more confident in this non-conflict version, one could assume that they can detect the conflict in the current version of the task, even if they do not manage to find the optimal strategy.

It would have been interesting to ask participants to estimate the quality of their selection phase as well as their overall performance. An optimal exploration is meant to help succeed

in the selection phase, thus leading to excellent performance. Therefore, final performance depends on the quality of both stages. By examining their confidence in both stages of the process, we could better understand whether participants distinguish between the two stages while recognising their complementarity. This would have allowed us to determine whether participants' confidence in their exploration influences their confidence in their final performance. If participants do not understand these connections, it is possible that their confidence in exploration is not linked to their confidence in performance.

The task's design is such that a poor exploration leaves little hope of achieving optimal performance. Understanding how participants perceive the relationship between the two stages could be important, especially to help them perform better in the task.

One potential avenue for future research would be to help participants enhance their strategy. An option could be to allow participants to perform only the selection stage, enabling them to realise how different exploration strategies can impact their performance. This would also help to assess whether participants are capable of making optimal selections following an optimal exploration.

As participants seem capable of choosing a strategy and being consistent in its application, teaching them an optimal strategy might lead them to apply it just as diligently. If they are encouraged to adopt a new strategy without explanation, it is likely they would prefer to stick with their current strategy, which gives them a good level of confidence.

## Bibliography

1. Ackerman, R., & Thompson, V. A. (2017). Meta-Reasoning: Monitoring and Control of Thinking and Reasoning. *Trends in Cognitive Sciences*, 21(8), 607–617. <https://doi.org/10.1016/j.tics.2017.05.004>
2. Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97. <https://doi.org/10.1016/j.jbef.2015.12.001>
3. De Neys, W. (2010). Heuristic Bias, Conflict, and Rationality in Decision-Making. In B. Glatzeder, V. Goel, & A. Müller (Eds.), *Towards a Theory of Thinking* (pp. 23–33). Springer Berlin Heidelberg. [https://doi.org/10.1007/978-3-642-03129-8\\_2](https://doi.org/10.1007/978-3-642-03129-8_2)
4. De Neys, W. (2012). Bias and Conflict: A Case for Logical Intuitions. *Perspectives on Psychological Science*, 7(1), 28–38. <https://doi.org/10.1177/1745691611429354>
5. De Neys, W., & Bonnefon, J.-F. (2013). The ‘whys’ and ‘whens’ of individual differences in thinking biases. *Trends in Cognitive Sciences*, 17(4), 172–178. <https://doi.org/10.1016/j.tics.2013.02.001>
6. De Neys, W., & Glumicic, T. (2008). Conflict monitoring in dual process theories of thinking. *Cognition*, 106(3), 1248–1299. <https://doi.org/10.1016/j.cognition.2007.06.002>
7. Duttler, K. (2016). Cognitive Skills and Confidence: Interrelations with Overestimation, Overplacement and Overprecision. *Bulletin of Economic Research*, 68(S1), 42–55. <https://doi.org/10.1111/boer.12069>
8. Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, 8. <https://doi.org/10.3389/fnhum.2014.00443>
9. Griffin, D., & Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24(3), 411–435. [https://doi.org/10.1016/0010-0285\(92\)90013-R](https://doi.org/10.1016/0010-0285(92)90013-R)
10. Jin, S., Verhaeghen, P., & Rahnev, D. (2022). Across-subject correlation between confidence and accuracy: A meta-analysis of the Confidence Database. *Psychonomic Bulletin & Review*, 29(4), 1405–1413. <https://doi.org/10.3758/s13423-022-02063-7>
11. Kassambara, A. (2023). *Ggpubr: 'ggplot2' Based Publication Ready Plots. R package version 0.6.0*. <https://CRAN.R-project.org/package=ggpubr>
12. Koriat, A. (2012). The self-consistency model of subjective confidence. *Psychological Review*, 119(1), 80–113. <https://doi.org/10.1037/a0025648>
13. Lichtenstein, S., Fischhoff, B., & Phillips, L. D. (1982). Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under Uncertainty* (1st ed., pp. 306–334). Cambridge University Press. <https://doi.org/10.1017/CBO9780511809477.023>
14. Lüdecke, D., S. Ben-Shachar, M., Patil, I., M. Wiernik, B., Bacher, E., Thériault, R., & Makowski, D. (2022). *easystats: Framework for Easy Statistical Modeling, Visualization, and Reporting*. CRAN. <https://easystats.github.io/easystats/>
15. Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2), 502–517. <https://doi.org/10.1037/0033-295X.115.2.502>
16. Nelson, T. O., & Narens, L. (1990). Metamemory: A Theoretical Framework and New Findings. In *Psychology of Learning and Motivation* (Vol. 26, pp. 125–173). Elsevier. [https://doi.org/10.1016/S0079-7421\(08\)60053-5](https://doi.org/10.1016/S0079-7421(08)60053-5)
17. Neys, W. D., Cromheeke, S., & Osman, M. (2011). Biased but in Doubt: Conflict and Decision Confidence. *PLOS ONE*, 6(1), e15954. <https://doi.org/10.1371/journal.pone.0015954>
18. O. Wilke, C. (2024). *Cowplot: Streamlined Plot Theme and Plot Annotations for 'ggplot2'*. R package version 1.1.3. <https://CRAN.R-project.org/package=cowplot>
19. R Core Team. (2023). *R: A Language and Environment for Statistical Computing*. <https://www.R-project.org/>
20. Šrol, J., & De Neys, W. (2021). Predicting individual differences in conflict detection and bias susceptibility during reasoning. *Thinking & Reasoning*, 27(1), 38–68. <https://doi.org/10.1080/13546783.2019.1708793>
21. Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., ...

### Chapter 3: Metacognition in set selection

Yutani, H. (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, 4(43), 1686.  
<https://doi.org/10.21105/joss.01686>

# Chapter 4: Learning and Adapting Optimal Strategies

Jean-Michel Dagba<sup>1</sup>, Jean-Christophe Vergnaud<sup>1,2</sup>, Vincent de Gardelle<sup>1,2,3</sup>

<sup>1</sup> Centre d'Économie de la Sorbonne (CES) & Université Paris 1 Panthéon Sorbonne, France

<sup>2</sup> Centre national de la recherche scientifique (CNRS), France

<sup>3</sup> Paris School of Economics (PSE), France

## Table of contents

Table of contents	115
I. Introduction	116
a. Leveraging Educational Interventions to Enhance Cognitive Skills and Bias Reduction	116
b. Financial incentives to improve effort	116
c. Combining Education and Incentives	117
d. Organising the educative content for exploration and selection stages	118
e. How do people generalise their (new) knowledge?	119
II. General method	119
a. Participants	119
b. Apparatus	120
c. Stimuli and procedure	120
d. Model and measures	126
e. Statistical analysis	130
III. Results	130
a. Main effects and interactions between Incentives and Education	130
b. Exploration stage	131
c. Selection stage	134
IV. Discussion	136
a. Study Objective and Main Findings	136
b. Effectiveness of Educational Interventions	137
c. Future Research Directions	137
Bibliography	139
Supplementary Materials	140

## I. Introduction

This study aims to evaluate the effectiveness of two forms of intervention—educational content and financial incentives—in improving decision-making performance and reducing exploration and selection biases in a set selection task. By examining how participants generalise their (newly acquired) skills to new versions of the task, we seek to understand the depth of their knowledge. Our approach offers new insights into how individuals navigate decision-making processes. Overall, we seek to uncover the cognitive mechanisms underlying complex multi-step reasoning.

### a. Leveraging Educational Interventions to Enhance Cognitive Skills and Bias Reduction

Educational interventions have proven effective in teaching individuals the principles of optimal decision-making strategies, with the potential for generalisation across tasks. A recent longitudinal study shows that educational debiasing interventions have positive effects that last in the long term (Morewedge et al., 2015). The authors demonstrated that educational interventions, such as games and videos, significantly reduced several cognitive biases, with effects lasting 8–12 weeks. Additionally, participants seemed able to use the intervention to reduce the effects of untreated forms of confirmation bias. Their findings provide a strong foundation for our use of educational content to address biases in decision-making, particularly in our set selection task. These interventions, therefore, seem quite effective and potentially generalisable to similar contexts. It seems appropriate to use educational content accompanied by a training phase with personalised feedback, drawing inspiration from the study by Morewedge (2015). There would be a theoretical section explaining an effective strategy for completing the two stages of the task, followed by a practical section to apply the theory with personalised feedback. We could then test whether participants manage to improve and generalise their new knowledge to solve other versions of the paradigm.

### b. Financial incentives to improve effort

In addition to educational intervention, we manipulate 3 types of financial incentives across experiments as another remediation factor. We conduct three experiments where financial incentives are manipulated to further motivate participants to maintain their attention and engage in completing the task.

Increasing financial incentives could make participants more inclined to consistently apply their best strategy. The behavioural economics literature has shown that financial incentives play a complex role in altering performance in experimental tasks. Camerer & Hogarth (1999) and (Voslinsky & Azar, 2021) conducted detailed reviews of the effects of financial incentives. Both reviews found that while such incentives generally increase cognitive effort in simple tasks, they do not always guarantee an improvement in performance, especially when tasks require high cognitive skills. In the same way, a study by Enke et al. (2023) shows that financial incentives have no effect on cognitive biases. Thus, it is unlikely that higher financial incentives will eliminate exploration and selection biases in our paradigm, but they may encourage participants to apply their best strategy more diligently and make fewer lapses in attention.

Research suggests that financial incentives involving potential losses (Gächter et al., 2022), can significantly impact participant motivation through loss aversion (Tversky & Kahneman, 1991). In this context, we hypothesise that incorporating losses alongside educational interventions may further enhance performance by encouraging participants to apply their strategies with greater rigour. These observations suggest that introducing an incentive system that includes potential losses could exert stronger behavioural pressure, leading to enhanced performance in tasks that require heightened attention.

### c. Combining Education and Incentives

To assess the relevance of educational content and financial incentives, we test their independent effects and a combination of the two factors. A study by Neumann et al. (2022) demonstrated that combining these factors improves participant performance. In their study, the authors conducted a series of experiments to analyse how education alone, financial incentives alone, and the combination of both influenced participants' use of decision-making rules. Their results showed that education played a central role in improving performance, but financial incentives enhanced participants' attention and motivation. The combination of the two interventions maximised the use of more optimal decision-making strategies, particularly by reducing errors related to inattention and encouraging greater rigour in the application of decision rules. Building on the work of Neumann et al. (2022), we apply a similar approach in our three experiments. By testing each factor independently and in combination, we aim to determine the most effective method for improving both attention and strategic decision-making in the context of our exploration-selection paradigm.

#### d. Organising the educative content for exploration and selection stages

Having established the theoretical framework for educational interventions and financial incentives, we now turn to the specific content and structure of our educative approach, focusing on the key stages of exploration and selection. Exploration and selection steps are link and need a common strategy to be well performed so our first objective is to resolve both the exploration and selection stages. Exploration stage consist in identifying the most useful information to make easier the selection stage which consist in choosing the higher total scores by comparing certain and uncertain scores.

Counterintuitively, it will be more pedagogically effective to explain to participants how to conduct useful exploration once they have understood how the information will be used in the selection stage. It is likely that participants struggle to develop a complex reasoning process that anticipates the set selection problem in order to determine how to conduct the exploration. Such challenges may stem from Failure of Contingent Thinking (FCT), where individuals, even when provided with all necessary information, fail to consider alternative states or hypothetical scenarios. As a result, people adopt strategies that overlook essential aspects of the task, leading to suboptimal decisions (Niederle & Vespa, 2023). In our paradigm, this is reflected in an exploration strategy focused on the options that appear most promising based on their score A, and a selection strategy that favours the explored options. Participants seem unable to identify the relevant scores to explore for a selection that is not limited to just the best option. In the selection phase, participants tend to undervalue the unknown scores B, favouring the known scores B instead. This bias leads to suboptimal selections, as they fail to properly account for the hidden information's potential value.

The selection educative content focused on interpreting the masked B scores. In the review on FCT by Esponda & Vespa (2023), it is shown that participants fail to adjust their decisions based on information that is not explicitly present but is implicit in the context. The selection educative content consisted of explaining that the values of the masked B scores can be estimated as 5s, allowing for easier comparisons of all options.

This then enabled us to explain the backward hypothetical reasoning involved for the exploration stage. As demonstrated by the study of Aina et al. (2023), individuals struggle to anticipate how their beliefs would change based on different potential contingencies. In their study, the authors investigated how individuals update their beliefs in response to new information across different contingent scenarios. They conducted a series of experiments

where participants were exposed to varying probabilities and potential outcomes, assessing how beliefs adjusted considering hypothetical situations. The results revealed that individuals often struggled with updating their beliefs accurately, particularly when faced with complex contingencies, leading to suboptimal decision-making. In our task, teaching the link between exploration and selection stages is precisely by describing the consequences of different explorations on the selection.

Exploration stage content is more extended because remediating the exploration seems to be the most consequential task, based on previous results. The first goal was then to demonstrate that exploring the favourites does not lead to an effective strategy for making a good selection. We expected to reduce significantly the bias of most participants and get their attention to explain how to realise the step. The content then introduces the key theoretical concept for conducting exploration, namely, exploring the information that will be most useful for making a set selection. After this theoretical definition, a practical example is provided to demonstrate how to identify the most useful options to explore.

#### e. How do people generalise their (new) knowledge?

A crucial objective of this study is to assess whether participants can generalise their newly acquired knowledge to different versions of the task. Effective generalisation is key to demonstrating the broader applicability of the educational content and strategies. By testing participants under increasingly challenging conditions, we aim to examine how they adjust their exploration and selection strategies. This phase of the study will offer valuable insights into the flexibility of the decision-making processes taught and the robustness of the educational content in fostering adaptive thinking.

## II. General method

### a. Participants

295 healthy French adults (50.5% of women) were included in this study. They were recruited from the Laboratoire d'Economie Expérimentale de Paris pool. They provided consent to take part in an experiment of 60 minutes.

In Experiment 8 (N = 100), participants earned an average of €8.10 during the tasks (in addition to a €5 payment for completing the study). For Experiment 9 (N = 87), they earned an average of €6.23 through the tasks (plus €9 for completing the experiment). For Experiment 10 (N = 108), they earned an average of €9.94 through the blocks (plus €8 for

completing the study). In the case of penalties due to performance during the task, participants were still guaranteed to earn €5.

After checking their answers to the three questions on understanding the probability rules used in the task, 13 participants were unfortunately withdrawn from the subsequent analyses because they had given at least 3 wrong answers. The analyses cover 282 participants (51.1% of women).

### b. Apparatus

Experiments were conducted online in July 2024. They were build using oTree (Chen et al., 2016) with the technical support of the Fédération S2CH.

### c. Stimuli and procedure

The main task used was the main task presented in chapters 1, 2 and 3. The task involves selecting the top 5 candidates from a pool of 10 based on two integer scores (score A and score B) ranging from 0 to 10, representing their abilities in skills A and B required for the job. The best candidates are those with the highest total score (A+B). Participants first have access to score A for each of the 10 candidates. They then choose 5 candidates from the 10 for which they want to see score B (exploration phase). Finally, they decide to recruit 5 candidates from the 10 (selection phase). Participants are informed that the scores are drawn from a uniform random distribution, with replacement. Score A and score B are independent and equally important. After the selection phase, participants receive feedback on the number of top 5 candidates they successfully identified. Candidates with equal scores are ranked equally. Therefore, several candidates might share ranks 1, 2, ..., 5. In such cases, all candidates ranked up to 5 are considered correct responses.

During the instructions, participants answer 3 multiple-choice questions to verify their understanding of the probability rules used in the task. After the instructions, participants complete 3 practice trials to familiarise themselves with the task.

During training, after the selection phase, participants receive richer feedback than in the actual trials. Here, they can see the unexplored score Bs and the complete ranking of the candidates. This helps them better understand how their performance is determined in each trial and recognise the existence of unexplored scores B.

After the intervention (described below), participants complete block 2, containing 10 trials with 5 explorations and 5 selections. After 2 blocks, in each group (Educated/Control), half of the participants complete block 3 (3 explorations – 6 selections) and then block 4 (3 explorations – 4 selections). The other half complete block 4 first, followed by block 3.

After the four blocks, participants complete a socio-demographic questionnaire and answer questions on the usefulness and clarity of the intervention they received.

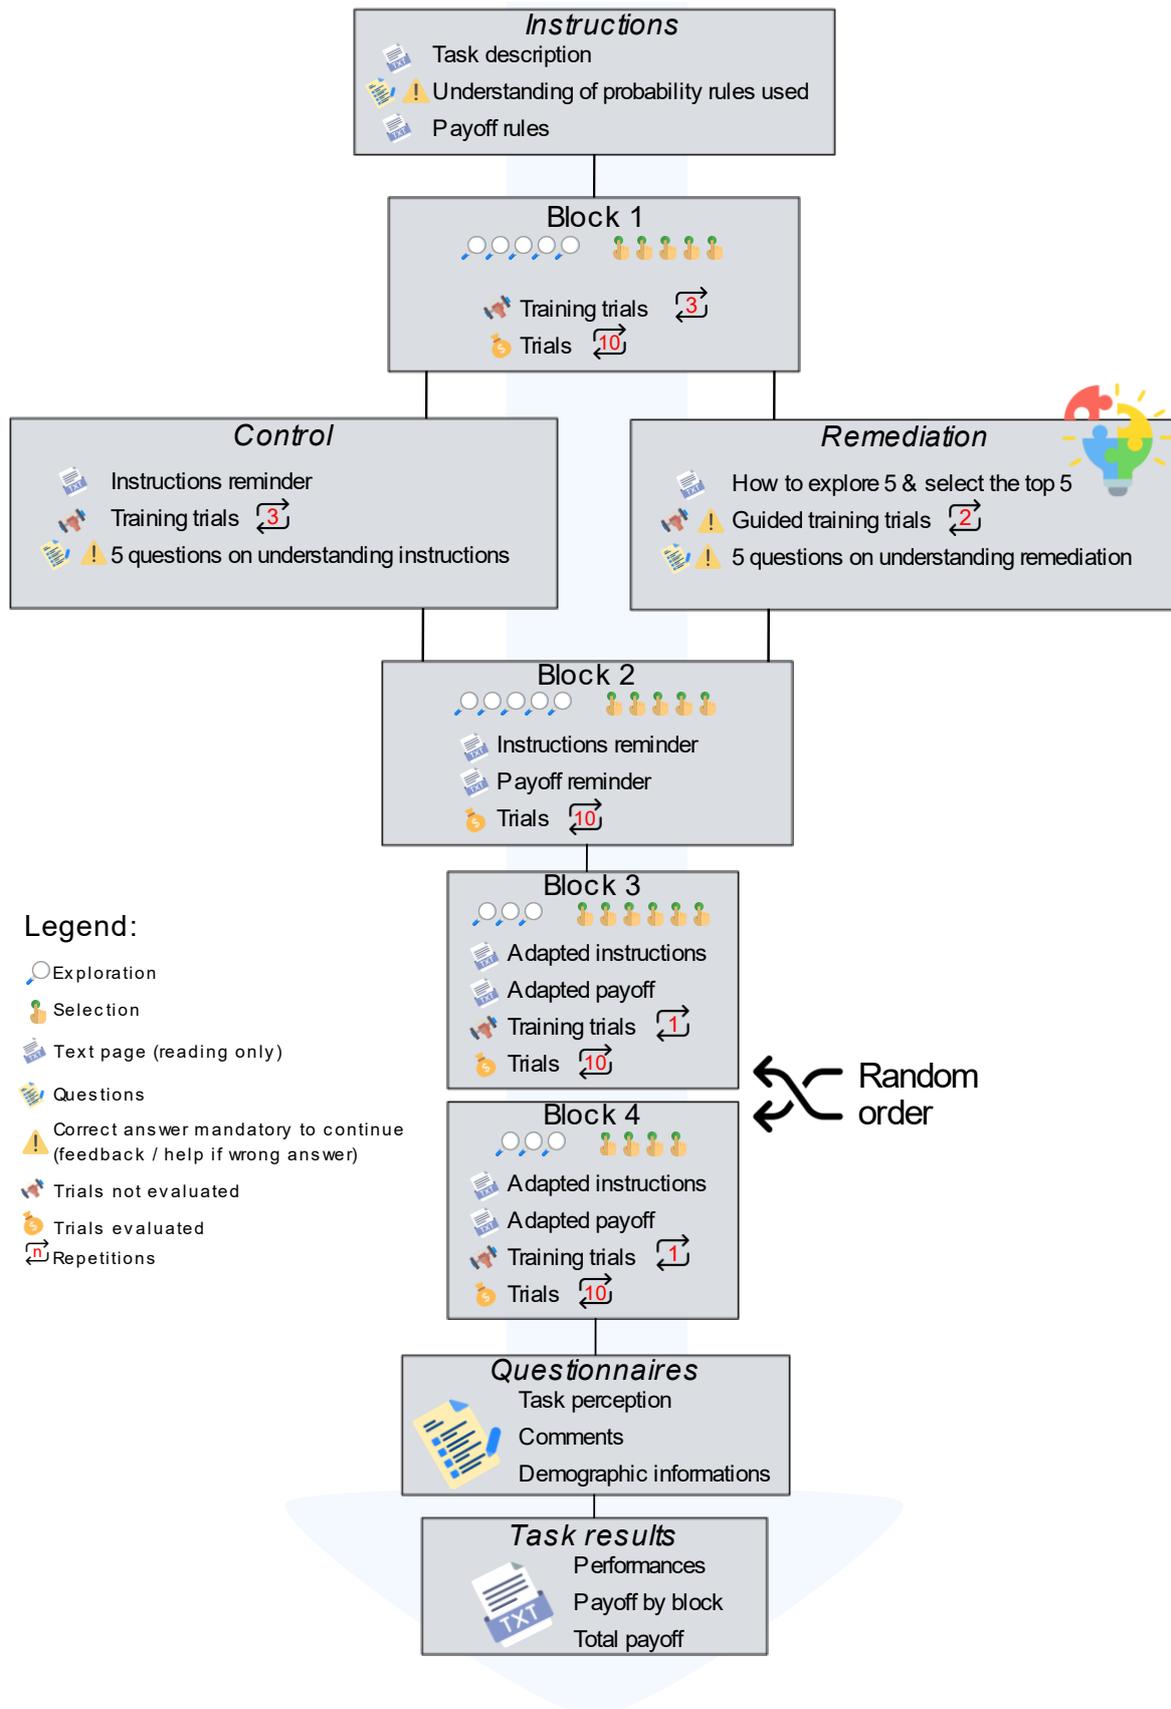


Fig 1.  
Task procedure

## Educative intervention

After block 1, participants undergo an intervention: half receive an educative content, and the other half a reminder of the instructions (Control group). In the Educated group, participants receive an illustrated presentation of a strategy for performing the exploration and selection steps with 5 explorations and 5 selections.

Our intervention aimed to help participants complete the two stages of the task in the same way as the algorithm we use to estimate the best solution for each trial. Our approach to the problem assumes that it is necessary to understand how to make the best possible selection to deduce which information is most useful to explore during the exploration stage. We therefore first present the content for selection, followed by the content for exploration (Fig 2).

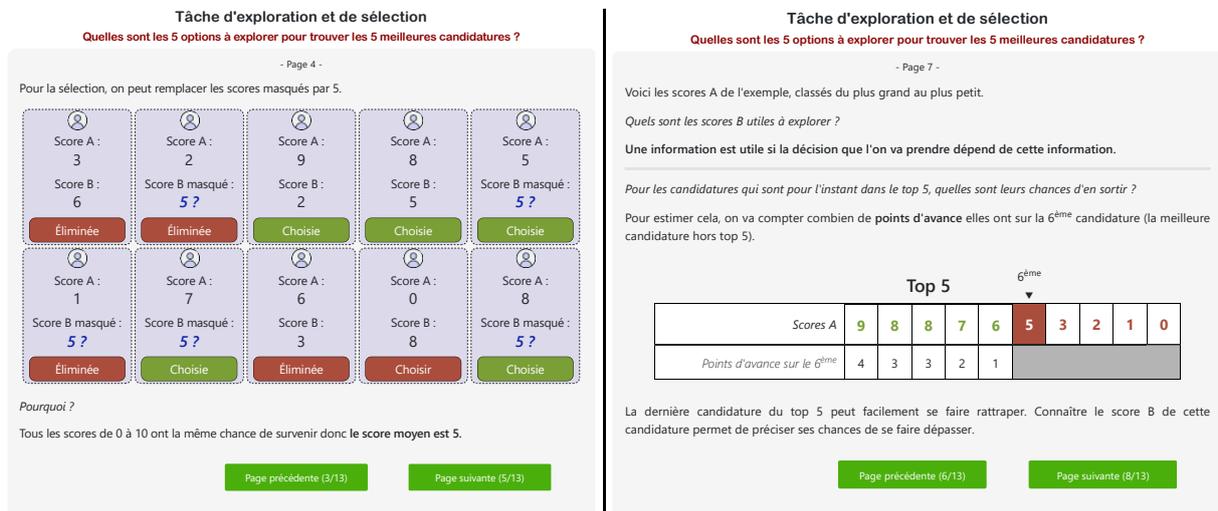


Fig 2.

Screenshots of two pages from the intervention received by the educated group. The left-hand image shows how to estimate the unknown B scores for selection (i.e., estimating them as scores of 5). The right-hand image explains how to identify the most useful candidates to explore from the provisional top 5, based on their lead over the provisionally 6th-ranked candidate. All screenshots can be found in the supplementary materials.

### Selection stage

For the selection phase, this strategy involves replacing unexplored score Bs with 5 (the average score). Then, the 5 candidates with the highest total scores (A+B) are selected without distinction between explored and unexplored candidates.

### *Exploration stage*

For the exploration phase, the strategy involves exploring candidates whose score B will be most useful in determining if they are in the top 5 or not. Two characteristics are considered: (1) the candidate's score A and (2) the candidate's provisional rank. The provisional rank is their position in the descending order of score A. These criteria estimate the likelihood of a candidate entering or exiting the top 5 given their provisional rank. We first calculate the number of points ahead of the first score A outside the top 5 (usually the 6th score A) for candidates in the provisional top 5. We calculate the number of points behind the last score A in the top 5 (usually the 5th score A) for candidates outside the provisional top 5. Candidates with the fewest points ahead (or behind) are most likely to lose their lead (or make up their deficit). These are the ones to prioritise for exploration. In case of a tie preventing the selection of exactly 5 candidates to explore, we then use the provisional ranking. We calculate the number of ranks ahead of the first rank outside the top 5 (usually the 6th rank) for candidates in the provisional top 5. We calculate the number of ranks behind the last rank in the top 5 (usually the 5th rank). Candidates with the fewest ranks ahead (or behind) are most likely to fall out of the top 5 (or enter it). Among the candidates that need to be distinguished, these are the ones to prioritise for exploration. If it is still not possible to identify exactly 5 candidates to explore after these two criteria, the strategy suggests breaking ties randomly.

### *Training with feedback after educative content*

After the presentation of this strategy, the participant must complete two guided training trials. They must correctly apply the strategy to proceed and receive feedback on each of their exploration and selection choices to aid them. Unlike the training trials conducted before block 1, there is no feedback on the selection result, as the objective here is solely to learn how to apply the presented strategy.

### *Understanding checks questions*

Finally, after the guided training, they answer five "True or False" questions to ensure their understanding of the presented strategy. If they make a mistake, they receive feedback explaining what they misunderstood. They must then correct their response to access the next question. The completion time for this educative content part is estimated at 10 minutes.

### **Intervention for control groups**

In Control groups, participants receive a reminder of the task instructions. It is specified that the reminder will be followed by three new training trials and questions on the instructions to ensure their proper understanding. The training trials function exactly the same way as

the training before block 1 (free exploration, free selection, feedback on the selection). After the training, the five "True or False" questions ensure the understanding of the instructions. As in the Educated group, participants receive feedback in case of an error and must correct themselves to proceed to the next question. The completion time for this Instructions Reminder part is estimated at 5-10 minutes.

### **Incentives manipulations**

We conducted three versions of this experiment to test different forms of financial incentives. As a reminder, each block allows for a monetary bonus. The bonus depends on the performance achieved in one of the 10 trials randomly selected for each block.

#### *Low Incentives (Exp. 8)*

In Exp. 8, each correct selection is worth +€0.50. For example, in Block 1, the maximum reward is relatively low (€2.50), and the difference between the potential gain with optimal performance and performance at chance level is quite small (less than €1).

#### *All-or-Nothing Incentives (Exp. 9)*

In Exp. 9, selecting 5 out of 5 correct candidates is worth +€5, while other outcomes (4/5, 3/5, 2/5, 1/5, 0/5) are worth €0. The maximum reward is therefore twice as high, but it is more difficult to achieve. There is thus a greater incentive to find the 5 best candidates in each trial for the block to potentially generate a gain. Conversely, performing at random has almost no chance of yielding a reward. The blocks operate on the same principle: a 5/5, 6/6, and 4/4 are required to earn €5, €6, and €4 in Blocks 2, 3, and 4, respectively.

#### *Bonus-Malus Incentives (Exp. 10)*

In Exp. 10, a correct selection is worth +€1 and an incorrect selection is worth -€2. Thus, in Block 1, selecting 4/5 or 5/5 correct candidates generates a potential gain, whereas selecting 0/5, 1/5, 2/5, or 3/5 correct candidates results in a potential loss. This time, performing the task at random almost always leads to a loss, as shown in Fig. 7 below. The other three blocks follow the same principle.

#### d. Model and measures

##### **Calculation of performance**

Performance for a given selection is defined as the fraction of selected options that verify the objective (top 5 in Blocks 1 and 2, top 6 in Block 3, top 4 in Block 4).

##### **Selections conditional on exploration**

As previously, we consider the optimal selection (i.e., the selection that maximises expected gains) based on known information (scores A and explored scores B). We calculated the optimal selection for the exploration conducted by the participants, named *Optimal\**. This selection estimates the best performance a participant can achieve given the scores B they chose to explore. We also calculated the optimal selection for optimal exploration (named *Optimal*), which is the best selection following the scores B revealed by optimal exploration (see below for the calculation of optimal exploration). We also calculated a random selection (to verify if participants perform the task correctly).

##### **Optimal exploration**

As in previous chapters, optimal exploration is defined as the exploration that maximises expected gain when followed by an optimal selection. We used the same approach to estimate optimal exploration from simulations. For each trial, an optimal exploration is estimated. We can then quantify the overlap between the participant's exploration in this trial and the optimal exploration. By averaging over all trials, we obtain a percentage representing the optimality of the participant's exploration.

##### **Suggested strategy by the educative intervention**

###### *Proxy for optimal exploration*

Participants received a simplified version of the steps used by our algorithm in the original task, with only 5 explorations and 5 selections (see intervention in supplementary materials). They were instructed to follow these steps:

1. Rank the candidates according to their A scores, from highest to lowest.
2. Identify the last score in the top 5 and the first score outside the top 5.

3. For candidates provisionally in the top 5: explore those with the smallest lead over the A score of the first candidate outside the top 5. For candidates provisionally outside the top 5: explore those with the smallest gap behind the A score of the last candidate in the top 5. Candidates with the smallest lead/gap are guaranteed to be explored.
4. In the event of tied candidates (equal lead or gap) preventing the exploration of exactly 5 candidates, another rule must be applied: count the rank differences in leads and gaps. A scores and their ranks may provide different information when there are point gaps between two consecutive candidates in the ranking. It is therefore possible to break ties using this characteristic. Then proceed as in point 3. For candidates provisionally in the top 5: explore those with the fewest rank differences ahead of the rank of the first candidate outside the top 5. For candidates provisionally outside the top 5: explore those with the fewest rank differences behind the rank of the last candidate in the top 5. This point 4 should only be applied to break ties when candidates are equal based on the previous lead/gap criterion.

As a reminder, due to ties, the last candidate in the top 5 is not necessarily ranked 5th. For example, they may be ranked 4th if another candidate has the same A score, and there are only 3 other candidates with better A scores.

5. If there are still ties preventing the identification of a set of 5 candidates to explore, participants are instructed to make a random choice as there is no further criterion to distinguish between tied candidates.

#### *Proxy for optimal selection*

As with exploration, participants received a simplified version of the steps used by our algorithm in the original task (5 explorations and 5 selections). They were instructed to follow these steps:

1. Replace all unknown scores B with 5. Each unknown score 8 could be any integer between 0 and 10. Therefore, the expected score is 5.
2. Select the candidates with the highest total score by summing the score A and the score B.
3. In the event of a tie between candidates preventing the completion of the top 5, participants are instructed to choose randomly between the tied candidates.

## Exploration strategies and efficiency categorisation

To analyse participants' exploration, we classify their strategies as *Efficient* or *Inefficient* for each block. An exploration is *Efficient* if it is closer to the optimal exploration (or the best heuristic) rather than to another predefined heuristic. For this, we used the estimate of optimal exploration (named *Optimal*) and the heuristics used in previous chapters. To recap, these heuristics are fixed strategies that explore options based on the ranking of A scores: exploring the highest-ranked options (*Highest* heuristic), the lowest-ranked options (*Lowest*), the middle-ranked options (*Middle*), and the extreme-ranked options (*Extreme*). We also consider random exploration (*Random*). These heuristics adapt to the number of options to explore but not to the number of options to select, making them more or less effective depending on the size of the selection set. To address this, we added the *Threshold* heuristic, which also adapts to the number of options to select: it explores options ranked around the selection threshold. For instance, when the objective is to find the top 5, the *Threshold* heuristic explores options with scores A ranked around the 5th and 6th positions. If the objective is to find the top 4, *Threshold* explores around the 4th and 5th positions. Note that for the top 5 case, *Threshold* overlaps with *Middle*, but for the top 4 case, *Threshold* shifts a little towards higher ranks. For the top 6 case, *Threshold* shifts a little towards lower ranks in the A score ranking. Thus, a participant is *Efficient* if they use an exploration strategy close to *Optimal* or *Threshold*. Participants classified as *Inefficient* follow another strategy.

### *Progress between blocks through learning and generalisation*

To evaluate the effectiveness of the educational intervention and instruction recall, we measured learning and generalisation using performance metrics based on participants' exploration strategies across multiple blocks in both the educated and control groups.

Learning was assessed by comparing participants' exploration strategies between Block 1 (before the intervention) and Block 2 (after the intervention). The proportion of participants who improved was defined as the part of participants who were categorised as *Inefficient* in Block 1 but became *Efficient* in Block 2, among all participants who were *Inefficient* in Block 1.

Generalisation was evaluated by examining participants' ability to transfer their newly acquired knowledge to novel variations of the task presented in Block 3 and Block 4. The generalisation measure focused on the proportion of participants who successfully adopted an *Efficient* strategy in these blocks compared to Block 1. As with learning, generalisation progress was measured by the proportion of participants who improved between Blocks 1 and 3 and between Blocks 1 and 4.

### Calculation of exploration and selection biases

As previously, the exploration bias towards favourites is the difference between the proportion of favourites selected by participants and the proportion of favourites selected by *Optimal\** (i.e., the optimal selection based on the participant's exploration). Favourites are the options with the highest scores A (the top 4 when the goal is to find the top 4, the top 5 when the goal is to find the top 5, etc.). A positive value indicates that participants explore too many options with the highest scores A. A negative value indicates that participants explore too many options with the lowest scores A.

Similarly, the selection bias towards explored options is the difference between the proportion of explored options selected by participants and the proportion of explored options selected by *Optimal\**. A positive value indicates a bias towards selecting explored options. A negative value indicates a bias towards selecting non-explored options.

We have added the calculation of another bias: the selection bias towards challengers. This is the difference between the proportion of non-favourite options selected by participants and the proportion of non-favourite options selected by *Optimal\**. A positive value indicates that participants select too many options with lower scores A (below the selection limit). A negative value indicates that participants select too many options with higher scores A.

### Calculation of mathematical errors

Finally, we calculated the *mathematical errors* made by a participant during their selection. This is done by adding the selection errors among explored options and those among unexplored options. Using the same proportions of explored and unexplored options as in the participants' selection, we simulate the best possible selection (named *Corrected*). If the participant's selection contains 3 explored options and 2 unexplored options, the *Corrected* selection choose the 3 options with the highest total scores among explored options and the 2 options with the highest scores A among unexplored options. We then count the number of differences between the participant's selection and the *Corrected* selection. Note that options with the same total scores among explored options (or the same scores A among unexplored options) are considered identical and not as differences. The number of differences between the two selections is divided by the number of options to be selected and expressed as a percentage. For example, a 40% mathematical error rate for a trial block where the top 5 must be found means that 2 out of 5 options were selected when a better score was visible.

### e. Statistical analysis

For pairwise comparisons of performance (between participants, or between models) we used Wilcoxon rank test (or Kruskal-Wallis rank test when there is more than 2 groups) and we report the effect size as  $r$  and 95% confidence intervals. We used chi-square tests and report Adjusted Cramer's  $v$  as the effect size to compare the proportions of the best-matching exploration strategy across experiments.

To measure the effects of the two forms of remediation across the tasks, we use a linear mixed model with an Intervention factor (Education vs. Control), an Incentives factor (Low vs. All-or-Nothing vs. Bonus-Malus), and a Block factor (Block 1 vs. Block 2 vs. Block 3 vs. Block 4). For the linear mixed model, we reported the analysis of Deviance (Type II Wald chi-square tests).

Statistical analyses were conducted using R (R Core Team, 2023), as well as the packages 'tidyverse' (Wickham et al., 2019), 'ggpubr' (Kassambara, 2023), 'cowplot' (O. Wilke, 2024), 'easystats' (Lüdtke et al., 2022), 'lmerTest' (Kuznetsova et al., 2017) and 'vcd' (Zeileis et al., 2007).

## III. Results

### a. Main effects and interactions between Incentives and Education

The study's analyses show that educational content influences the exploration and selection stages. Financial incentives have no effect on the exploration stage, and there is only a main effect on the selection bias. The following sections present the results for the two stages of the task in more detail.

#### Understanding of the Educational Intervention

Participants generally appeared to have understood the educational content. 95% of participants made fewer than 3 errors out of 5 verification questions (with 62% making no errors at all). For participants who had a reminder of the instructions, 100% made fewer than 3 errors on the 5 questions (with 77% making no errors). After an error, participants received an explanatory reminder and had to correct their response before continuing with the task.

## b. Exploration stage

The exploration stage is clearly impacted by the educational intervention provided to half of the participants. However, it does not vary at all based on the different financial incentives. Since there is no interaction between Intervention and Incentives, the results below combine the three experiments that involved different financial incentive systems. The results thus present the 282 participants divided into two groups: educated (N = 132) and control (N = 150).

### **Exploration optimality**

Regarding the optimality of participants' exploration, the linear mixed model indicates main effects of the Intervention (Chisq = 69.1229, Df = 1,  $p < .001$ ) and the Block (Chisq = 448.9040, Df = 3,  $p < .001$ ), as well as an Intervention x Block interaction (Chisq = 111.4959, Df = 3,  $p < .001$ ). As shown in Fig. 3a & 3b, the educated group demonstrated more optimal exploration than the control group after the interventions.

### **Exploration bias**

One of our main findings in previous experiments was the presence of an exploration bias towards favoured options, i.e., options with the highest scores and thus the most likely to be among the final selection set. The linear mixed model shows a main effect of Block (Chisq = 143.1994, Df = 3,  $p < .001$ ), and a significant Intervention x Block interaction (Chisq = 30.7831, Df = 3,  $p < .001$ ).

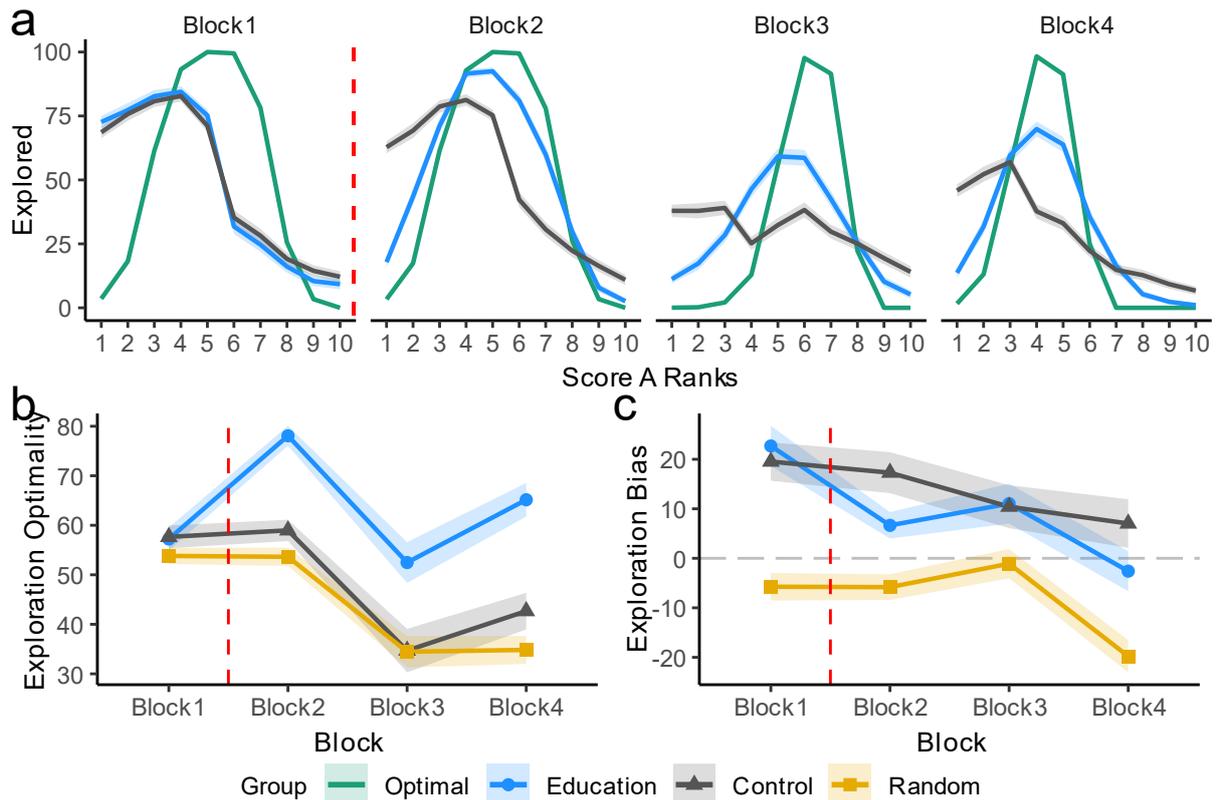


Fig. 3.

Panel (a) represents the options explored by participants and by the optimal exploration strategy according to the ranks of the options' scores A, for each block. Panel (b) shows the optimality of exploration by participants in the Remediation group (in blue) and the Control group (in grey). The optimality of a random exploration is represented (in yellow). Panel (c) illustrates the exploration bias towards favoured options for the Remediation group (in blue) and the Control group (in grey) across the blocks. The exploration bias of a random exploration is represented (in yellow). The dashed redlines represent the intervention.

### Exploration strategies

The proportion of participants using an efficient strategy varies significantly by Block and Intervention ( $\chi^2 = 212.50, p < .001; v = .43$ ), as shown in detail by block in Fig. 4a. In the Control groups, the proportion of participants using an efficient exploration strategy (optimal or near optimal) remains stable across the blocks ( $\chi^2 = 1.81, p = .612$ ), whereas it varies significantly among the educated participants ( $\chi^2 = 129.30, p < .001; v = .49$ ).

### Learning

There are almost 8 times more participants who improved their exploration strategy between Blocks 1 and 2 in the educated group compared to the control group ( $\chi^2 = 105.71, p < .001; v = .69, 95\% \text{ CI} = [.58, 1.00]$ ). Indeed, there is an equivalent proportion between the groups of participants who were just as efficient before and after the intervention (~18%). Moreover, among participants who were not using an efficient strategy before the intervention, 79% improved due to the educational intervention, while only 11% improved due to the instruction

recall. Overall, there is a significant and large difference in knowledge after the intervention between the groups ( $\chi^2 = 108.53$ ,  $p < .001$ ;  $v = .62$ , 95% CI = [.51, 1.00]), as shown in Fig. 4b.

### *Generalisation*

Educated participants were also more likely to generalise their knowledge to solve variations of the task than participants in the control condition. Between Blocks 1 and 3, nearly twice as many participants successfully switched to an efficient strategy ( $\chi^2 = 6.51$ ,  $p = .011$ ; Adjusted Cramer's  $v = .16$ , 95% CI = [.02, 1.00]). Between Blocks 1 and 4, nearly four times as many participants progressed in the educated group compared to the control group ( $\chi^2 = 51.59$ ,  $p < .001$ ; Adjusted Cramer's  $v = .47$ , 95% CI = [.36, 1.00]).

In more detail, participants (from the educated group) better generalised the intervention in Block 4 (3 explorations – 4 selections) than in Block 3 (3 explorations – 6 selections). As shown in Fig. 3a, the optimal strategy shifts according to the selection size. When the selection size increases, as in Block 3, exploration shifts towards ranks 6 to 10. When the selection size decreases, optimal exploration shifts towards ranks 1 to 5. Thus, the distribution of optimal exploration across ranks varies significantly between blocks ( $\chi^2 = 89.41$ ,  $p < .001$ ; Adjusted Cramer's  $v = .54$ , 95% CI = [.44, 1.00]). The distribution of exploration performed by educated participants across option ranks varies significantly between Blocks 2, 3, and 4 ( $\chi^2 = 15.67$ ,  $p < .001$ ; Adjusted Cramer's  $v = .21$ , 95% CI = [.09, 1.00]). This difference in exploration distribution is not significant between Blocks 2 and 3 ( $\chi^2 = 1.76$ ,  $p = .185$ ). Participants did not significantly adjust their strategy when moving from Block 2 (5 explorations & 5 selections) to Block 3 (3 explorations & 6 selections). The difference is significant between Blocks 2 and 4 ( $\chi^2 = 5.65$ ,  $p = .017$ ; Adjusted Cramer's  $v = .16$ , 95% CI = [.00, 1.00]), indicating that they successfully adapted their exploration strategy to complete Block 4 (3 explorations & 4 selections).

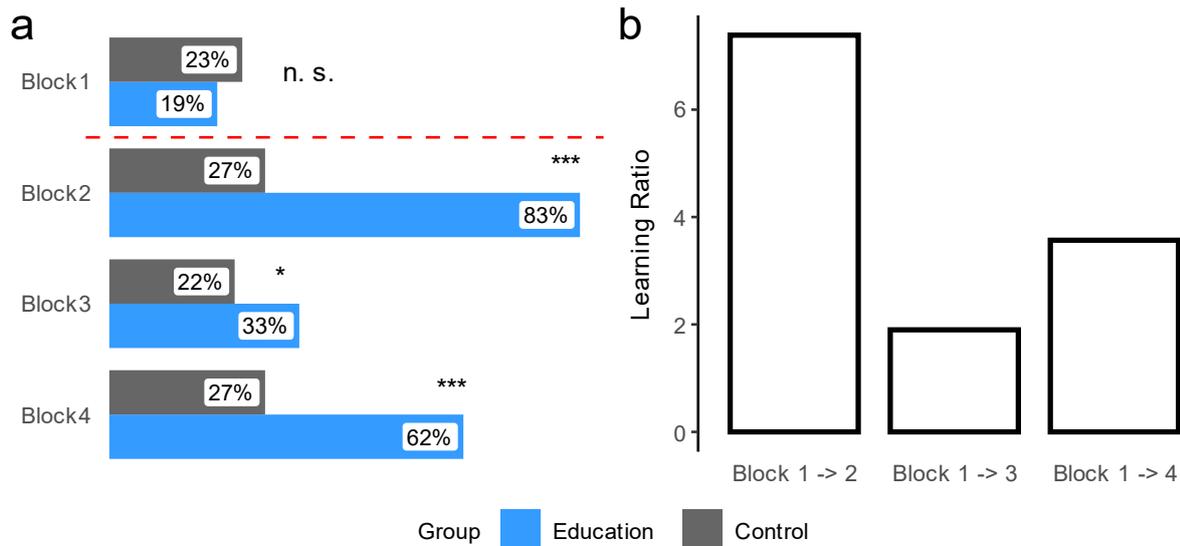


Fig. 4.

Panel (a) presents the percentage of participants who used an effective strategy according to the block and the type of intervention received. Education = intervention after Block 1 explaining the optimal strategy; Control = intervention after Block 1 reminding the instructions. Panel (b) presents the ratio between the learning of educated and uneducated participants for blocks 2, 3 and 4 compared with block 1. The ratio is the percentage of participants in the educated group who succeeded in adopting an effective exploration strategy after intervention divided by the percentage of participants in the control group who succeeded in adopting an effective exploration strategy after intervention. Participants who already had an effective exploration strategy before and after the intervention were not counted.

### c. Selection stage

The selection stage is also mainly influenced by the educational content, except the selection bias. Therefore, the results presented below intervention group participants regardless of the financial incentives they received (except for the selection bias). Overall, educated participants chose better and performed better than participants in control groups.

### Selection optimality and Performance

Selection optimality varies significantly according to the intervention ( $\chi^2 = 13.1777$ ,  $Df = 1$ ,  $p < .001$ ), the block ( $\chi^2 = 162.0236$ ,  $Df = 3$ ,  $p < .001$ ), and the Intervention x Block interaction ( $\chi^2 = 14.0422$ ,  $Df = 3$ ,  $p = .003$ ). As shown in Fig. 5, the effect size between groups is larger in Block 2 ( $r = .35$ ) than in Block 3 ( $r = .23$ ) and in Block 4 ( $r = .22$ ).

As a result, there are also main effects of the intervention ( $\chi^2 = 27.3574$ ,  $Df = 1$ ,  $p < .001$ ) and the Block ( $\chi^2 = 70.8123$ ,  $Df = 3$ ,  $p < .001$ ), and an interaction effect of Block x Group ( $\chi^2 = 24.5081$ ,  $Df = 3$ ,  $p < .001$ ) on final performance. There is no other significant effect on performance.

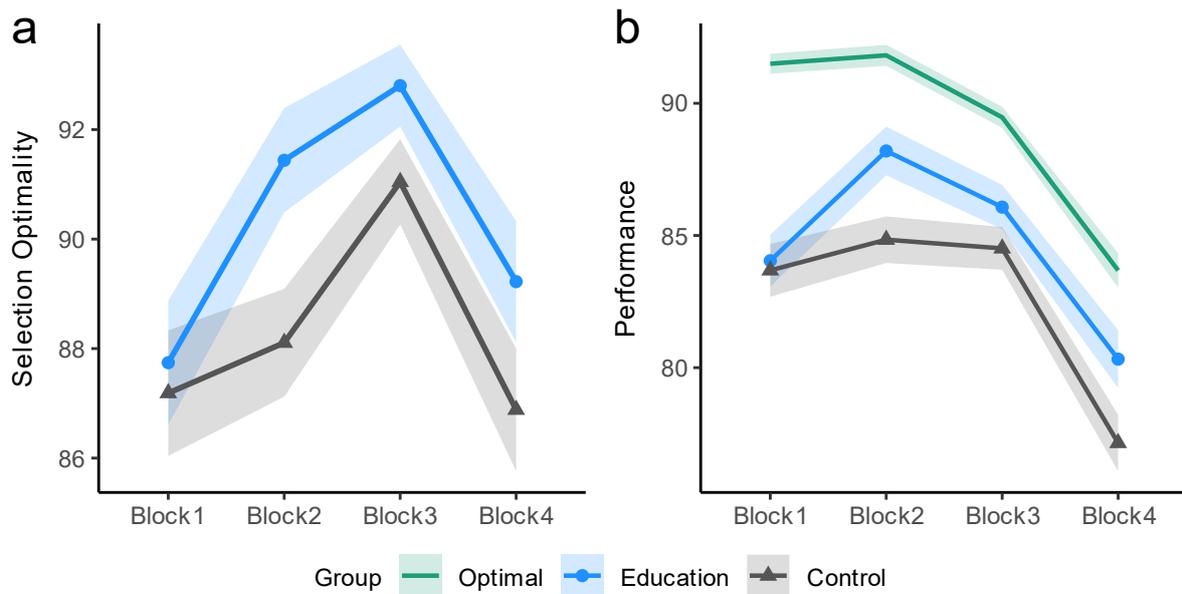


Fig. 5.

Graphs illustrating the effects of interventions on the selection stage, according to the block. Panel (a) presents the percentage of correspondence between the Optimal\* selection and the selection made by the participants. Panel (b) represents final performance, which is the proportion of selected options that were indeed among the best. Green line represents an optimal agent performance.

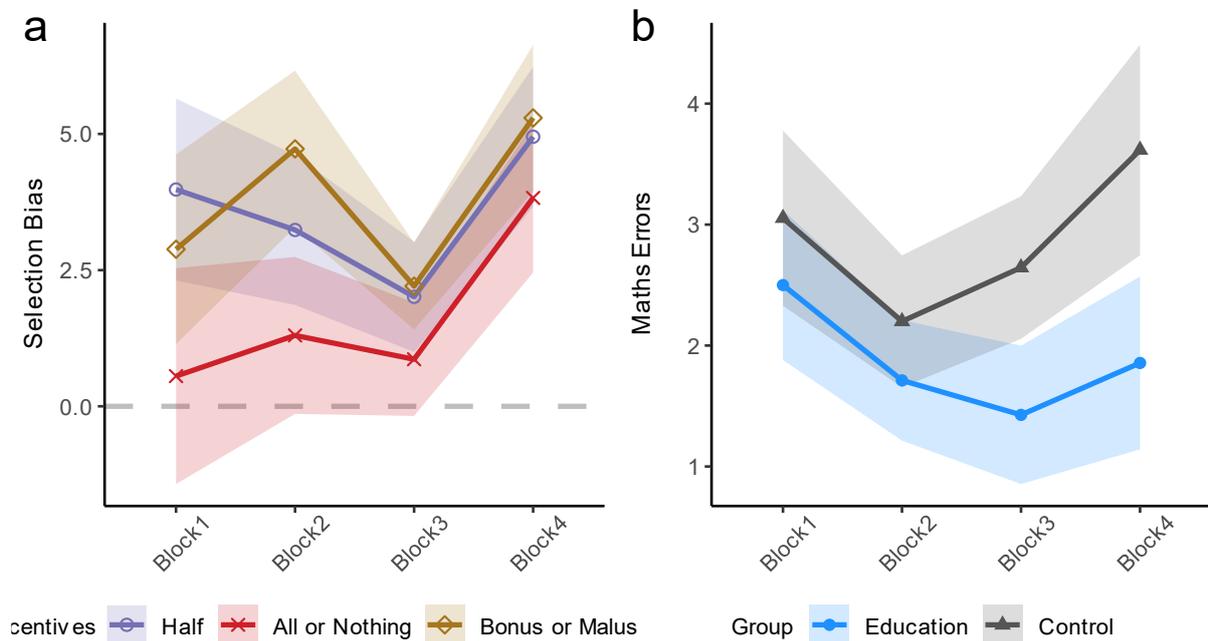
### Selection bias and mathematical errors

Selection bias towards explored options seem to be affected by financial incentives and maths errors are affected by the intervention. Although the effects are significant, it is also observed that the selection bias is rather small (3 percentage points between Optimal and participants), as is the proportion of mathematical errors (less than 3% errors). This yields results that are difficult to interpret beyond their description.

There is a main effect of Financial Incentives ( $\text{Chisq} = 9.3353$ ,  $\text{Df} = 2$ ,  $p = .009$ ) and a main effect of the Block ( $\text{Chisq} = 48.6568$ ,  $\text{Df} = 3$ ,  $p < .001$ ) on selection bias. There is no effect of the intervention ( $\text{Chisq} = .0068$ ,  $\text{Df} = 1$ ,  $p = .934$ ), nor any significant interaction effect. Participants in the All or Nothing version (Exp. 9) had significantly less selection bias than those in the Half version ( $W = 52742.00$ ,  $p < .001$ ;  $r = -.18$ , 95% CI =  $[-.26, -.10]$ ) and the Bonus or Malus version ( $W = 56200.50$ ,  $p < .001$ ;  $r = -.17$ , 95% CI =  $[-.25, -.09]$ ). There was no difference between the Half and Bonus or Malus versions ( $W = 79859.50$ ,  $p = .827$ ).

Mathematical errors vary significantly according to the Intervention ( $\text{Chisq} = 6.9984$ ,  $\text{Df} = 1$ ,  $p = .008$ ), the Block ( $\text{Chisq} = 22.4365$ ,  $\text{Df} = 3$ ,  $p < .001$ ), the Block x Intervention interaction ( $\text{Chisq} = 10.3333$ ,  $\text{Df} = 3$ ,  $p = .016$ ), and the three-way Block x Intervention x Incentives interaction ( $\text{Chisq} = 14.5229$ ,  $\text{Df} = 6$ ,  $p = .024$ ). The double interaction should be interpreted with caution as the number of errors is very low and could simply be due to chance. Fig. 6

presents these variables based on their main effects: selection bias according to each financial incentive format and mathematical errors according to the Intervention, for each block.



**Fig. 6.** Panel (a) presents the selection bias towards the explored options. This is the difference between the proportion of explored options selected by the participants and the proportion of explored options selected by the Optimal\* (i.e. the best possible selection given the participants' exploration). Panel (b) presents the proportion of options selected by participants that were dominated by an option for which they had equivalent information.

## IV. Discussion

### a. Study Objective and Main Findings

The primary goal of this study was to assess the effectiveness of two forms of intervention—educational content and financial incentives—on improving information search and decision-making performance in a set selection task. We specifically aimed to reduce exploration and selection biases and to test the ability of participants to generalise their newly acquired skills to more complex versions of the task. By investigating how individuals engage in multi-step reasoning, this study contributes to understanding the cognitive processes involved in decision-making under uncertainty.

Our results provide clear evidence of the effectiveness of educational interventions in improving exploration and selection strategies. Participants who received educational content showed a marked improvement in exploration strategies, with 8 times more educated participants than not-educated participants adopting an optimal strategy thanks to the intervention.

Furthermore, the educated participants are able to partially generalise their new knowledge to solve variations of the task. They have a slightly more effective exploration strategy and better performance than participants in the control group in the configuration with 3 explorations and 6 selections. They show much better exploration and improved performance compared to the control group participants in the version with 3 explorations and 4 selections. The exploration and selection of participants are not affected by financial incentives. We observe only an effect on selection bias, reflecting a lower bias in Exp. 9 (All or Nothing).

### b. Effectiveness of Educational Interventions

The findings related to exploration strategies align well with prior research on contingent thinking and decision-making biases (Esponda & Vespa, 2023; Martínez-Marquina et al., 2019; Niederle & Vespa, 2023). Participants often fail to consider alternative scenarios or hypothetical outcomes, resulting in suboptimal decisions. Our educational intervention successfully addressed these cognitive limitations by providing explicit strategies for exploration, which improved participants' exploration even in modified versions of the task. This supports the idea that educational interventions, particularly when coupled with feedback, can lead to great improvements in cognitive reasoning (Morewedge et al., 2015).

The lack of a significant effect of financial incentives highlights the complexity of this cognitive task. Financial incentives may be more effective for simpler tasks that require less cognitive load, as suggested by Camerer & Hogarth (1999). Our results are consistent with studies that show that they do not necessarily lead to a better reasoning (Gächter et al., 2022).

For the selection phase, our findings suggest that financial incentives involving penalties for incorrect decisions may not trigger loss aversion, prompting participants to apply more careful and deliberate strategies. This contradicts Tversky & Kahneman (1991), who demonstrated that loss aversion can have a profound impact on decision-making. The similar performance in the Bonus-Malus condition (Exp 10) compared to Exp. 8 and 9 indicates that penalties for errors did not sharpen participants' focus and effort.

### c. Future Research Directions

While we observed significant improvements in exploration and selection strategies for a large proportion of participants, a subset of participants did not benefit from the remediation. It remains unclear whether this is due to differences in cognitive abilities,

motivation, or understanding of the task. At the same time, some participants managed to solve the original task and its variations perfectly, even without the educational intervention. A closer examination of this subset of participants could reveal whether their unique perception of the task enables more accurate reasoning. Sharing their approach could refine the educational intervention and offer greater support to those who continue to struggle despite the remediation.

Continuing research on the quality of learning at the metacognitive level could help us better understand what changed for participants after the intervention and during the generalisation phase. To do this, we could ask participants to provide metacognitive judgments on how closely their strategy after the intervention aligns with the taught strategy. One might then question whether the ability to judge one's own learning predicts performance not only in the original task but also in variations of the original task. In addition to refining the educational intervention, it is important to explore how well participants can generalise their learning to new contexts and tasks.

Furthermore, the generalisation aspect could be fully investigated. Some participants achieved it brilliantly, while others struggled. This difficulty may arise from challenges in extracting a general rule from specific cases or from confusion about the broader applicability of the remediation strategies presented. To help participants generalise the remediation without ambiguity, it would be useful to: (1) clearly draw their attention to the need to adapt the remediation to solve tasks with different parameters, (2) explain how the number of explorations and selections influences the strategy to be used, and (3) offer training with feedback on several versions of the task. In this way, participants would have a greater chance of understanding the overall strategy, regardless of the number of explorations and selections in the task. We might then ask ourselves what types of real-life tasks could be better performed through learning this complex reasoning.

In conclusion, this study underscores the vital role of tailored educational interventions in enhancing complex decision-making processes. The need for a cognitive approach to improve reasoning is critical. Future research should continue to explore how these interventions can be refined, combined and generalised to other cognitive tasks, with the goal of improving decision-making in both experimental and real-world settings.

## Bibliography

1. Aina, C., Amelio, A., & Brütt, K. (2023). *Contingent belief updating* (Working Paper 263). ECONtribute Discussion Paper. <https://www.econstor.eu/handle/10419/283305>
2. Camerer, C. F., & Hogarth, R. M. (1999). The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. *Journal of Risk and Uncertainty*, 19(1), 7–42. <https://doi.org/10.1023/A:1007850605129>
3. Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97. <https://doi.org/10.1016/j.jbef.2015.12.001>
4. Enke, B., Gneezy, U., Hall, B., Martin, D., Nelidov, V., Offerman, T., & van de Ven, J. (2023). Cognitive Biases: Mistakes or Missing Stakes? *The Review of Economics and Statistics*, 105(4), 818–832. [https://doi.org/10.1162/rest\\_a\\_01093](https://doi.org/10.1162/rest_a_01093)
5. Esponda, I., & Vespa, E. (2023). Contingent Thinking and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory. *The Review of Economic Studies*, rdad102. <https://doi.org/10.1093/restud/rdad102>
6. Gächter, S., Johnson, E. J., & Herrmann, A. (2022). Individual-level loss aversion in riskless and risky choices. *Theory and Decision*, 92(3), 599–624. <https://doi.org/10.1007/s11238-021-09839-8>
7. Kassambara, A. (2023). *Ggpubr: 'ggplot2' Based Publication Ready Plots. R package version 0.6.0*. <https://CRAN.R-project.org/package=ggpubr>
8. Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). **lmerTest** Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82(13). <https://doi.org/10.18637/jss.v082.i13>
9. Lüdtke, D., S. Ben-Shachar, M., Patil, I., M. Wiernik, B., Bacher, E., Thériault, R., & Makowski, D. (2022). *easystats: Framework for Easy Statistical Modeling, Visualization, and Reporting*. CRAN. <https://easystats.github.io/easystats/>
10. Martínez-Marquina, A., Niederle, M., & Vespa, E. (2019). Failures in Contingent Reasoning: The Role of Uncertainty. *American Economic Review*, 109(10), 3437–3474. <https://doi.org/10.1257/aer.20171764>
11. Morewedge, C. K., Yoon, H., Scopelliti, I., Symborski, C. W., Korris, J. H., & Kassam, K. S. (2015). Debiasing Decisions: Improved Decision Making With a Single Training Intervention. *Policy Insights from the Behavioral and Brain Sciences*, 2(1), 129–140. <https://doi.org/10.1177/2372732215600886>
12. Neumann, M., Hengeveld, M., Niessen, A. S. M., Tendeiro, J. N., & Meijer, R. R. (2022). Education increases decision-rule use: An investigation of education and incentives to improve decision making. *Journal of Experimental Psychology: Applied*, 28(1), 166–178. <https://doi.org/10.1037/xap0000372>
13. Niederle, M., & Vespa, E. (2023). Cognitive Limitations: Failures of Contingent Thinking. *Annual Review of Economics*, 15(Volume 15, 2023), 307–328. <https://doi.org/10.1146/annurev-economics-091622-124733>
14. O. Wilke, C. (2024). *Cowplot: Streamlined Plot Theme and Plot Annotations for 'ggplot2'*. *R package version 1.1.3*. <https://CRAN.R-project.org/package=cowplot>
15. R Core Team. (2023). *R: A Language and Environment for Statistical Computing*. <https://www.R-project.org/>
16. Tversky, A., & Kahneman, D. (1991). Loss Aversion in Riskless Choice: A Reference-Dependent Model\*. *The Quarterly Journal of Economics*, 106(4), 1039–1061. <https://doi.org/10.2307/2937956>
17. Voslinsky, A., & Azar, O. H. (2021). Incentives in experimental economics. *Journal of Behavioral and Experimental Economics*, 93, 101706. <https://doi.org/10.1016/j.socec.2021.101706>
18. Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., ... Yutani, H. (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, 4(43), 1686. <https://doi.org/10.21105/joss.01686>
19. Zeileis, A., Meyer, D., & Hornik, K. (2007). Residual-Based Shadings for Visualizing (Conditional) Independence. *Journal of Computational and Graphical Statistics*, 16(3), 507–525. <https://doi.org/10.1198/106186007X237856>

## Supplementary materials

### Replication of the previous main results

The average performance level is 83% (95% CI = [79.52%, 83.31%]), which corresponds to an average normalised performance of 89% (95% CI = [86.66%, 90.92%]). This performance level is consistent across experiments (Kruskal-Wallis  $\chi^2 = 4.41$ ,  $p = .221$ ) and between the two conditions in Exp. 8 ( $W = 1858.00$ ,  $p = 0.133$ ). On average, the performance loss during the exploration stage is 4.69 percentage points (95% CI = [4.29, 5.08]), and the performance loss during the selection stage is 4.37 percentage points (95% CI = [3.56, 5.18]). These losses are not significantly different from those observed in previous experiments for exploration (Kruskal-Wallis  $\chi^2 = 3.16$ ,  $p = .368$ ) or selection (Kruskal-Wallis  $\chi^2 = 7.57$ ,  $p = .056$ ). There is also no difference between the Remediation group and the Control group in terms of loss during exploration ( $W = 1584.00$ ,  $p = .947$ ) or selection ( $W = 1356.50$ ,  $p = .168$ ).

Categorising participants based on the strategy they are closest to a majority are again classified as Highest (62%). The categorisation of participants by strategies (Highest, Lowest, Middle, Extreme, Random, and Optimal) does not change across experiments ( $\chi^2 = 21.45$ ,  $p = 0.123$ ) nor between the Remediation and Recall groups in Exp. 8 ( $\chi^2 = 3.65$ ,  $p = .601$ ).

Furthermore, participants' exploration corresponds to 58% of the optimal exploration, which is similar to the overlap observed in previous versions of the experiment (Kruskal-Wallis  $\chi^2 = 1.82$ ,  $p = 0.612$ ) and consistent between the Remediation and Recall groups in Exp. 8 (Kruskal-Wallis  $\chi^2 = 1.75$ ,  $p = .185$ ).

The exploration bias towards favourites is significantly positive ( $W = 5018.00$ ,  $p < .001$ ;  $r = 0.71$ , 95% CI = [.58, .80]), similar to previous experiments (Kruskal-Wallis  $\chi^2 = 2.91$ ,  $p = .406$ ), and there is no difference between the Recall and Remediation groups ( $W = 1740.00$ ,  $p = 0.409$ ).

In the selection stage, the two types of errors identified previously are also observed. On the one hand, the bias of selection towards explored options is significantly positive on average ( $W = 4773.00$ ,  $p < .001$ ;  $r = 0.65$ , 95% CI [0.51, 0.76]), as in previous experiments (Kruskal-Wallis  $\chi^2 = 5.22$ ,  $p = .156$ ), and there is no difference between the two treatment groups ( $W = 1735.00$ ,  $p = .424$ ). On the other hand, mathematical errors result in a significant performance loss of 2.95 (95% CI = [1.67, 4.22]) percentage points on average ( $W = 17596.00$ ,  $p < .001$ ;  $r = 0.60$ , 95% CI = [.53, 1.00]), as in previous experiments (Kruskal-Wallis  $\chi^2 = 2.05$ ,  $p = .562$ ), with no difference between treatment groups ( $W = 1682.00$ ,  $p = .612$ ).

## Screenshots of Instructions

### Lancement de l'expérience

#### Description générale de l'expérience

Vous allez prendre part à une expérience dont le temps de complétion est estimé entre 60 et 70 minutes.

Cette estimation temporelle est uniquement indicative afin de vous aider à décider si vous pouvez réaliser cette expérience dès maintenant et en une seule fois.

Pendant l'expérience, nous vous demandons de vous consacrer entièrement à l'étude afin de nous fournir des résultats fiables et utilisables pour la recherche scientifique. Merci de mettre vos notifications en silencieux et de vous installer dans un endroit calme.

L'expérience comporte 5 parties. Les consignes spécifiques à chaque partie vous seront précisées le moment venu.

Pour obtenir la gratification dont il est question pour certaines parties, vous devez compléter l'intégralité de l'étude. Un gain de 5,00€ sera ajouté à vos gains cumulés dans les différentes parties.

Si besoin, vous pouvez quitter l'étude définitivement à tout moment.

Si un bug survient pendant l'expérience (écran bloqué, bouton qui ne répond plus, etc.), veuillez fermer votre navigateur, puis l'ouvrir de nouveau et revenir sur l'expérience avec le lien qui vous a été fourni (votre progression devrait être restaurée).

Vous pouvez appuyer sur "Page suivante" dès que vous voulez commencer. Merci de votre participation.

Page suivante

### Instructions - Page 1 / 7

#### Description générale de la partie 1

Vous allez prendre part à une tâche dont l'objectif est de réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Lors de chaque processus, votre mission est de **sélectionner les 5 meilleures candidatures parmi une liste de 10 candidatures**.

Pour chaque candidature, vous aurez accès à **un score A allant de 0 (inclus) à 10 (inclus)** correspondant à une évaluation d'une compétence A de chaque personne.

Parmi les 10 candidatures, vous pourrez alors indiquer **5 candidatures** pour lesquelles vous aurez accès à **un score B** correspondant à une évaluation d'une compétence B de chaque personne.

Enfin, vous devrez prendre une décision de recrutement, à savoir **choisir selon vous les 5 meilleures candidatures** parmi les 10 candidatures.

Selon nos estimations basées sur les tests de l'expérience, le temps nécessaire pour réaliser cette partie est de **10 à 15 min** (temps de lecture des instructions compris).

Page suivante

## Instructions - Page 1 / 7

### Description générale de la partie 1

Vous allez prendre part à une tâche dont l'objectif est de réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Lors de chaque processus, votre mission est de **sélectionner les 5 meilleures candidatures parmi une liste de 10 candidatures**.

Pour chaque candidature, vous aurez accès à **un score A allant de 0 (inclus) à 10 (inclus)** correspondant à une évaluation d'une compétence A de chaque personne.

Parmi les 10 candidatures, vous pourrez alors indiquer **5 candidatures** pour lesquelles vous aurez accès à **un score B** correspondant à une évaluation d'une compétence B de chaque personne.

Enfin, vous devrez prendre une décision de recrutement, à savoir **choisir selon vous les 5 meilleures candidatures** parmi les 10 candidatures.

Vous pouvez obtenir des gains durant cette partie en fonction de la qualité **d'un de vos 10 processus de recrutement tiré au hasard**. Si, pour ce processus de recrutement tiré au hasard, vous avez trouvé les **5** meilleures candidatures, votre gain sera de **5,00€**. Si vous avez trouvé 0, 1, 2, 3 ou 4 candidatures parmi les meilleures, vous n'obtiendrez **aucun gain (0,00€) pour cette partie**.

Selon nos estimations basées sur les tests de l'expérience, le temps nécessaire pour réaliser cette partie est de **10 à 15 min** (temps de lecture des instructions compris).

Page suivante

## Instructions - Page 1 / 7

### Description générale de la partie 1

Vous allez prendre part à une tâche dont l'objectif est de réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Lors de chaque processus, votre mission est de **sélectionner les 5 meilleures candidatures parmi une liste de 10 candidatures**.

Pour chaque candidature, vous aurez accès à **un score A allant de 0 (inclus) à 10 (inclus)** correspondant à une évaluation d'une compétence A de chaque personne.

Parmi les 10 candidatures, vous pourrez alors indiquer **5 candidatures** pour lesquelles vous aurez accès à **un score B** correspondant à une évaluation d'une compétence B de chaque personne.

Enfin, vous devrez prendre une décision de recrutement, à savoir **choisir selon vous les 5 meilleures candidatures** parmi les 10 candidatures.

Vous pouvez avoir un bonus ou un malus sur votre gain durant cette partie en fonction de la qualité **d'un de vos 10 processus de recrutement tiré au hasard**. Pour le processus qui sera tiré au hasard, une bonne recrue vous rapportera 1,00€ et une mauvaise recrue vous fera perdre 2,00€ :

- 5 / 5 candidatures sélectionnées parmi les meilleures → Bonus de 5,00€ ( +5 € )
- 4 / 5 candidatures sélectionnées parmi les meilleures → Bonus de 2,00€ ( +2 € )
- 3 / 5 candidatures sélectionnées parmi les meilleures → Malus de 1,00€ ( -1 € )
- 2 / 5 candidatures sélectionnées parmi les meilleures → Malus de 4,00€ ( -4 € )
- 1 / 5 candidatures sélectionnées parmi les meilleures → Malus de 7,00€ ( -7 € )
- 0 / 5 candidatures sélectionnées parmi les meilleures → Malus de 10,00€ ( -10 € )

Selon nos estimations basées sur les tests de l'expérience, le temps nécessaire pour réaliser cette partie est de **10 à 15 min** (temps de lecture des instructions compris).

Page suivante

## Instructions - Page 2 / 7

### Première étape - Exploration

Lors de la première étape (l'exploration), vous verrez apparaître simultanément à l'écran les 10 candidatures ainsi qu'un score A allant de 0 à 10, correspondant à une évaluation de la compétence A. Plus le score A est élevé, plus la personne a été évaluée comme possédant la compétence A, nécessaire pour le poste en question.

Les candidatures seront présentées de la manière suivante :

 Score A : <b>5</b> Score B : Voir score B	 Score A : <b>7</b> Score B : Voir score B	 Score A : <b>3</b> Score B : Voir score B	 Score A : <b>3</b> Score B : Voir score B	 Score A : <b>3</b> Score B : Voir score B
 Score A : <b>4</b> Score B : Voir score B	 Score A : <b>10</b> Score B : Voir score B	 Score A : <b>7</b> Score B : Voir score B	 Score A : <b>9</b> Score B : Voir score B	 Score A : <b>9</b> Score B : Voir score B

Pour chaque candidature, vous pouvez cliquer sur le bouton **"Voir score B"** si vous souhaitez **l'explorer**, c'est-à-dire **voir le score d'évaluation de la compétence B de la personne correspondante** lors de l'étape suivante.

Vous devez indiquer **exactement 5 candidatures à explorer** : vous ne pourrez passer à l'étape suivante que lorsque vous aurez indiqué 5 personnes dont vous voulez voir le score B.

Pour chaque candidature, **vous pouvez changer d'avis** autant de fois que vous le souhaitez en cliquant de nouveau sur le même bouton.

Si vous avez déjà indiqué 5 candidatures à explorer, vous devez en retirer une avant de pouvoir en ajouter une autre.

Le nombre de clics que vous effectuez ainsi que vos temps de réponse sont enregistrés, mais cela n'a pas d'influence sur vos gains potentiels.

Vos demandes d'exploration sont enregistrées et définitives lorsque vous cliquez sur le bouton "Valider mes demandes d'exploration et passer à l'étape de sélection", en bas de page.

Page suivante

### Instructions - Page 3 / 7

#### Deuxième étape - Sélection

Lors de la deuxième étape (la sélection), vous verrez les 10 mêmes candidatures avec le score A que vous aviez déjà vu précédemment et **le score B pour les 5 candidatures que vous avez demandé à explorer lors de l'étape précédente**. Pour les autres personnes, le score B est remplacé par un " ? " pour indiquer qu'il existe aussi, mais que vous n'en avez pas connaissance. Plus le score B est élevé, plus la personne a été évaluée comme possédant la compétence B, nécessaire pour le poste en question.

Les candidatures seront présentées de la manière suivante :

 Score A : <b>5</b> Score B : <b>7</b> <input type="button" value="Choisir"/>	 Score A : <b>7</b> Score B : <b>?</b> <input type="button" value="Choisir"/>	 Score A : <b>3</b> Score B : <b>3</b> <input type="button" value="Choisir"/>	 Score A : <b>3</b> Score B : <b>?</b> <input type="button" value="Choisir"/>	 Score A : <b>3</b> Score B : <b>3</b> <input type="button" value="Choisir"/>
 Score A : <b>4</b> Score B : <b>?</b> <input type="button" value="Choisir"/>	 Score A : <b>10</b> Score B : <b>1</b> <input type="button" value="Choisir"/>	 Score A : <b>7</b> Score B : <b>?</b> <input type="button" value="Choisir"/>	 Score A : <b>9</b> Score B : <b>2</b> <input type="button" value="Choisir"/>	 Score A : <b>9</b> Score B : <b>?</b> <input type="button" value="Choisir"/>

L'étape fonctionne de la même manière que l'étape précédente.

Pour chaque candidature, vous pouvez cliquer sur le bouton "**Choisir**" si vous souhaitez recruter cette personne.

**Votre objectif est de sélectionner les 5 candidatures qui ont les meilleurs scores en faisant l'addition de leurs deux scores (A et B). Pour déterminer les meilleures candidatures, les deux scores de chaque personne sont pris en compte, que vous ayez décidé d'explorer le score B ou non.**

Vous devez sélectionner **exactement 5 candidatures** : vous ne pourrez valider votre recrutement que lorsque vous aurez recruté 5 personnes.

**Toutes les candidatures peuvent être choisies** : vous pouvez choisir les candidatures pour lesquelles vous avez vu un seul score comme les candidatures pour lesquelles vous avez vu les deux scores.

Pour chaque candidature, **vous pouvez changer d'avis** autant de fois que vous le souhaitez en cliquant de nouveau sur le même bouton.

Si vous avez déjà indiqué 5 candidatures à sélectionner, vous devez en retirer une avant de pouvoir en choisir une autre.

Le nombre de clics que vous effectuez ainsi que vos temps de réponse sont enregistrés, mais cela n'a pas d'influence sur vos gains potentiels.

Vos décisions de recrutement sont enregistrées et définitives lorsque vous cliquez sur le bouton "Achever le processus et enregistrer mes recrutements", en bas de page.

Page suivante

## Instructions - Page 4 / 7

### Résultat du recrutement

À la suite de votre recrutement, vous aurez une indication sur le nombre de candidatures choisies qui étaient effectivement parmi les meilleures candidatures.

Le résultat sera présenté sous la forme suivante :

Parmi les **5** candidatures sélectionnées, **4** étaient parmi les meilleures candidatures.

Pour vous permettre de mieux comprendre comment est calculé le résultat, voici une vue plus complète à laquelle vous aurez accès uniquement pendant l'entraînement :

Parmi les **5** candidatures sélectionnées, **4** étaient parmi les meilleures candidatures.

**Les candidatures sont classées dans l'ordre décroissant selon la somme de leurs deux scores. La candidature ayant le score total le plus élevé est première et la candidature ayant le score total le plus faible est dernière.**

Toutes les candidatures classées entre la 1<sup>ère</sup> et la 5<sup>ème</sup> place sont parmi les meilleures candidatures (soit 5 candidatures ou plus en cas d'égalité).

Les candidatures classées au-delà de la 5<sup>ème</sup> place étaient les moins bonnes candidatures, c'est-à-dire celles qu'il fallait éliminer.

 3 <sup>ème</sup> Score A : <b>5</b> Score B : <b>7</b> Choisie - Bonne décision	 6 <sup>ème</sup> Score A : <b>7</b> Score B : <b>3</b> Choisie - Mauvaise décision	 8 <sup>ème</sup> Score A : <b>3</b> Score B : <b>3</b> Éliminée	 6 <sup>ème</sup> Score A : <b>3</b> Score B : <b>7</b> Éliminée	 8 <sup>ème</sup> Score A : <b>3</b> Score B : <b>3</b> Éliminée
 10 <sup>ème</sup> Score A : <b>4</b> Score B : <b>0</b> Éliminée	 4 <sup>ème</sup> Score A : <b>10</b> Score B : <b>1</b> Choisie - Bonne décision	 1 <sup>ère</sup> Score A : <b>7</b> Score B : <b>8</b> Éliminée	 4 <sup>ème</sup> Score A : <b>9</b> Score B : <b>2</b> Choisie - Bonne décision	 2 <sup>ème</sup> Score A : <b>9</b> Score B : <b>5</b> Choisie - Bonne décision

Parmi les 5 candidatures sélectionnées, choisir la deuxième candidature de la rangée du haut était une mauvaise décision, car elle est classée à la 6<sup>ème</sup> place, c'est-à-dire après la **5<sup>ème</sup> place**.

À l'inverse, la candidature située au milieu de la rangée du bas était 1<sup>ère</sup>, mais n'a pas été choisie.

En cas d'égalité, les candidatures sont considérées ex-æquo et sont classées à la même place. Si elles sont entre la 1<sup>ère</sup> et la 5<sup>ème</sup> place, elles sont considérées comme des bonnes réponses **indifféremment**.

Pendant l'expérience, le résultat de vos décisions de recrutement est **affiché pendant 5 secondes** avant que le processus de recrutement suivant ne commence.

### Instructions - Page 5 / 7

#### Calcul des scores

Les scores A et B des candidatures sont des **nombre entiers tirés aléatoirement** de 0 à 10. Ils ont tous la **même chance** d'être tirés au sort.

Ainsi, chaque score est **indépendant** et a la **même probabilité de survenir** quelque soit les autres scores. En d'autres termes, tous les scores A sont indépendants, tous les scores B sont indépendants, et les scores A et B de chaque candidature sont indépendants.

Avant de continuer, merci de répondre aux 3 questions ci-dessous. Un message en vert s'affichera lorsque vous aurez validé la réponse correcte. Vous pourrez passer à la page suivante lorsque toutes les réponses validées seront correctes.

#### Question 1

La probabilité d'avoir un score de 7 est-elle supérieure, inférieure ou égale à la probabilité d'avoir un score de 1 ?

- Supérieure à la probabilité d'avoir un score de 1
- Inférieure à la probabilité d'avoir un score de 1
- Égale à la probabilité d'avoir un score de 1

Valider

Les deux prochaines questions seront basées sur l'exemple de deux candidatures : Camille possédant un score A de 8 et Charlie possédant un score A de 3.

	
Camille	Charlie
Score A :	Score A :
8	3
Score B :	Score B :

#### Question 2

La probabilité que le score B de Camille soit supérieur au score B de Charlie est-elle supérieure à 50%, inférieure à 50% ou égale à 50% ?

- Supérieure à 50%
- Inférieure à 50%
- Égale à 50%

Valider

#### Question 3

Un score B qui serait égal à 3 est-il plus probable pour Camille, plus probable pour Charlie, ou également probable pour Camille et Charlie ?

- Plus probable pour Camille
- Plus probable pour Charlie
- Également probable

Valider

Page suivante

**Question 1**

La probabilité d'avoir un score de 7 est-elle supérieure, inférieure ou égale à la probabilité d'avoir un score de 1 ?

- Supérieure à la probabilité d'avoir un score de 1
- Inférieure à la probabilité d'avoir un score de 1
- Égale à la probabilité d'avoir un score de 1

Mauvaise réponse. Pour rappel, chaque nombre a exactement la même probabilité d'être tiré.

Valider

Les deux prochaines questions seront basées sur l'exemple de deux candidatures : Camille possédant un score A de 8 et Charlie possédant un score A de 3.

 Camille	 Charlie
Score A : 8	Score A : 3
Score B :	Score B :

**Question 2**

La probabilité que le score B de Camille soit supérieur au score B de Charlie est-elle supérieure à 50%, inférieure à 50% ou égale à 50% ?

- Supérieure à 50%
- Inférieure à 50%
- Égale à 50%

Mauvaise réponse. Pour rappel, le score A et B sont strictement indépendants (avoir 10 ou 0 au score A n'influence pas le score B).

Valider

**Question 3**

Un score B qui serait égal à 3 est-il plus probable pour Camille, plus probable pour Charlie, ou également probable pour Camille et Charlie ?

- Plus probable pour Camille
- Plus probable pour Charlie
- Également probable

Mauvaise réponse. Pour rappel, les valeurs déjà visibles n'influencent pas celle non découverte.

Valider

**Question 1**

La probabilité d'avoir un score de 7 est-elle supérieure, inférieure ou égale à la probabilité d'avoir un score de 1 ?

- Supérieure à la probabilité d'avoir un score de 1
- Inférieure à la probabilité d'avoir un score de 1
- Égale à la probabilité d'avoir un score de 1

Bravo ! Les nombres ayant tous la même chance de survenir, la probabilité d'obtenir un 7 est similaire à celle d'obtenir un 1.

Valider

Les deux prochaines questions seront basées sur l'exemple de deux candidatures : Camille possédant un score A de 8 et Charlie possédant un score A de 3.

 Camille	 Charlie
Score A : 8	Score A : 3
Score B :	Score B :

**Question 2**

La probabilité que le score B de Camille soit supérieur au score B de Charlie est-elle supérieure à 50%, inférieure à 50% ou égale à 50% ?

- Supérieure à 50%
- Inférieure à 50%
- Égale à 50%

Bravo ! Comme le score A et le score B sont indépendants, le score A n'influence pas le score B.

Valider

**Question 3**

Un score B qui serait égal à 3 est-il plus probable pour Camille, plus probable pour Charlie, ou également probable pour Camille et Charlie ?

- Plus probable pour Camille
- Plus probable pour Charlie
- Également probable

Bravo ! Comme pour la deuxième question, tous les scores sont indépendants.

Valider

## Instructions - Page 6 / 7

### Résumé

Vous allez réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Chaque processus de recrutement considère **10 candidatures** pour le poste.

Pour chaque processus, **votre objectif est de sélectionner les 5 meilleures candidatures.**

1) Étape d'exploration : Après avoir consulté un score d'évaluation d'une compétence A pour toutes les candidatures, vous devez **indiquer 5 candidatures pour lesquelles vous souhaitez voir un score d'évaluation d'une compétence B.**

Chaque score est compris entre **0 (inclus) et 10 (inclus)** et plus il est **élevé, plus** la personne est évaluée comme étant **compétente.**

2) Étape de sélection : Vous devez alors **recruter les 5 personnes qui sont au total les plus compétentes (score A + score B).**

**Toutes les candidatures peuvent être choisies** : vous pouvez choisir les candidatures pour lesquelles vous avez vu un seul score comme les candidatures pour lesquelles vous avez vu les deux scores.

Tous les scores A sont indépendants, tous les scores B sont indépendants, et les scores A et B pour chaque candidature sont indépendants.

Vous obtiendrez des gains pour la réalisation de la tâche et en fonction de la qualité **d'un de vos recrutements tiré au hasard.** Chaque candidature choisie qui est effectivement parmi les meilleures candidatures vous rapporte **0,50€**, soit un gain total maximum de **2,50€.**

Si vous avez terminé toutes les parties de l'expérience, vous recevrez vos gains quelques jours après la fin de l'expérience.

Le temps de réalisation de cette partie de l'expérience est estimé à **10 à 15 minutes** (instructions comprises).

Page suivante

## Instructions - Page 6 / 7

### Résumé

Vous allez réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Chaque processus de recrutement considère **10 candidatures** pour le poste.

Pour chaque processus, **votre objectif est de sélectionner les 5 meilleures candidatures**.

1) Étape d'exploration : Après avoir consulté un score d'évaluation d'une compétence A pour toutes les candidatures, vous devez **indiquer 5 candidatures pour lesquelles vous souhaitez voir un score d'évaluation d'une compétence B**.

Chaque score est compris entre **0 (inclus) et 10 (inclus)** et plus il est **élevé, plus** la personne est évaluée comme étant **compétente**.

2) Étape de sélection : Vous devez alors **recruter les 5 personnes qui sont au total les plus compétentes (score A + score B)**.

**Toutes les candidatures peuvent être choisies** : vous pouvez choisir les candidatures pour lesquelles vous avez vu un seul score comme les candidatures pour lesquelles vous avez vu les deux scores.

Tous les scores A sont indépendants, tous les scores B sont indépendants, et les scores A et B pour chaque candidature sont indépendants.

Vous pourrez obtenir un **gain de 5,00€** si vous avez choisi les **5** meilleures candidatures pour l'un de vos 10 processus de recrutement **déterminé au hasard**. Vous n'obtiendrez pas de gain dans cette partie si votre sélection n'était pas parfaite durant le processus tiré au hasard.

Le temps de réalisation de cette partie de l'expérience est estimé à **10 à 15 minutes** (instructions comprises).

Page suivante

## Instructions - Page 6 / 7

### Résumé

Vous allez réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Chaque processus de recrutement considère **10 candidatures** pour le poste.

Pour chaque processus, **votre objectif est de sélectionner les 5 meilleures candidatures.**

1) Étape d'exploration : Après avoir consulté un score d'évaluation d'une compétence A pour toutes les candidatures, vous devez **indiquer 5 candidatures pour lesquelles vous souhaitez voir un score d'évaluation d'une compétence B.**

Chaque score est compris entre **0 (inclus) et 10 (inclus)** et plus il est **élevé, plus** la personne est évaluée comme étant **compétente.**

2) Étape de sélection : Vous devez alors **recruter les 5 personnes qui sont au total les plus compétentes (score A + score B).**

**Toutes les candidatures peuvent être choisies** : vous pouvez choisir les candidatures pour lesquelles vous avez vu un seul score comme les candidatures pour lesquelles vous avez vu les deux scores.

Tous les scores A sont indépendants, tous les scores B sont indépendants, et les scores A et B pour chaque candidature sont indépendants.

Vous pouvez avoir un bonus ou un malus sur votre gain durant cette partie en fonction de la qualité **d'un de vos 10 processus de recrutement tiré au hasard.** Pour le processus qui sera tiré au hasard :

- 5 / 5 → Bonus de 5,00€ ( +5€ )
- 4 / 5 → Bonus de 2,00€ ( +2€ )
- 3 / 5 → Malus de 1,00€ ( -1€ )
- 2 / 5 → Malus de 4,00€ ( -4€ )
- 1 / 5 → Malus de 7,00€ ( -7€ )
- 0 / 5 → Malus de 10,00€ ( -10€ )

Le temps de réalisation de cette partie de l'expérience est estimé à **10 à 15 minutes** (instructions comprises).

Page suivante

## Instructions - Page 7 / 7

### Entraînement

Avant de commencer l'expérience, vous avez **3 essais d'entraînement** afin de vous familiariser avec la tâche. **Vos réponses ne seront pas comptabilisées.** L'expérience commence **immédiatement** après l'entraînement.

Appuyez sur « Commencer » pour lancer l'entraînement.

Commencer

## Screenshots of Training trials

**Essai d'entraînement n°1 / 1 - étape 1 : exploration**  
**Explorer 3 candidatures puis sélectionner les 4 meilleures parmi les 10**

Score A : 2 Score B : ? Score B caché	Score A : 5 Score B : ? Score B à voir	Score A : 9 Score B : ? Score B caché	Score A : 0 Score B : ? Score B caché	Score A : 2 Score B : ? Score B caché
Score A : 1 Score B : ? Score B caché	Score A : 6 Score B : ? Score B à voir	Score A : 5 Score B : ? Score B caché	Score A : 9 Score B : ? Score B caché	Score A : 4 Score B : ? Score B à voir

Valider mes demandes d'exploration et passer à l'étape de sélection

**Essai d'entraînement n°1 / 1 - étape 2 : sélection**  
**Explorer 3 candidatures puis sélectionner les 4 meilleures parmi les 10**

Score A : 2 Score B : ? Éliminée	Score A : 5 Score B : 6 Choisie	Score A : 9 Score B : ? Choisie	Score A : 0 Score B : ? Éliminée	Score A : 2 Score B : ? Éliminée
Score A : 1 Score B : ? Éliminée	Score A : 6 Score B : ? Éliminée	Score A : 5 Score B : 0 Éliminée	Score A : 9 Score B : ? Choisie	Score A : 4 Score B : 9 Choisie

Achever le processus et enregistrer mes choix

**Essai d'entraînement n°1 / 1 - résultats**  
 Parmi les 4 candidatures sélectionnées, 3 sont parmi les meilleures candidatures.  
 Toutes les candidatures classées entre la 1<sup>ère</sup> et la 4<sup>ème</sup> place sont parmi les meilleures candidatures (soit 4 candidatures ou plus en cas d'égalité).  
 Les candidatures classées au-delà de la 4<sup>ème</sup> place étaient finalement les moins bonnes.

1 <sup>ère</sup> Score A : 2 Score B : 3 Éliminée	2 <sup>ème</sup> Score A : 5 Score B : 6 Choisie - Bonne décision	3 <sup>ème</sup> Score A : 9 Score B : 0 Choisie - Mauvaise décision	4 <sup>ème</sup> Score A : 0 Score B : 10 Éliminée	5 <sup>ème</sup> Score A : 2 Score B : 3 Éliminée
6 <sup>ème</sup> Score A : 1 Score B : 4 Éliminée	7 <sup>ème</sup> Score A : 6 Score B : 5 Éliminée	8 <sup>ème</sup> Score A : 5 Score B : 0 Éliminée	9 <sup>ème</sup> Score A : 9 Score B : 5 Choisie - Bonne décision	10 <sup>ème</sup> Score A : 4 Score B : 9 Choisie - Bonne décision

Continuer

**Essai d'entraînement n°1 / 1 - étape 1 : exploration**  
**Explorer 3 candidatures puis sélectionner les 6 meilleures parmi les 10**

Score A : 9 Score B : ? Score B caché	Score A : 7 Score B : ? Score B caché	Score A : 2 Score B : ? Score B caché	Score A : 1 Score B : ? Score B caché	Score A : 3 Score B : ? Score B à voir
Score A : 0 Score B : ? Score B caché	Score A : 3 Score B : ? Score B à voir	Score A : 3 Score B : ? Score B à voir	Score A : 1 Score B : ? Score B caché	Score A : 3 Score B : ? Score B caché

Valider mes demandes d'exploration et passer à l'étape de sélection

**Essai d'entraînement n°1 / 1 - étape 2 : sélection**  
**Explorer 3 candidatures puis sélectionner les 6 meilleures parmi les 10**

Score A : 9 Score B : ? Choisie	Score A : 7 Score B : ? Choisie	Score A : 2 Score B : ? Choisie	Score A : 1 Score B : ? Éliminée	Score A : 3 Score B : 5 Choisie
Score A : 0 Score B : ? Éliminée	Score A : 3 Score B : 0 Éliminée	Score A : 3 Score B : 9 Choisie	Score A : 1 Score B : ? Éliminée	Score A : 3 Score B : ? Choisie

Achever le processus et enregistrer mes choix

**Essai d'entraînement n°1 / 1 - résultats**  
 Parmi les 6 candidatures sélectionnées, 6 sont parmi les meilleures candidatures.  
 Toutes les candidatures classées entre la 1<sup>ère</sup> et la 6<sup>ème</sup> place sont parmi les meilleures candidatures (soit 6 candidatures ou plus en cas d'égalité).  
 Les candidatures classées au-delà de la 6<sup>ème</sup> place étaient finalement les moins bonnes.

1 <sup>ère</sup> Score A : 9 Score B : 1 Choisie - Bonne décision	2 <sup>ème</sup> Score A : 7 Score B : 3 Choisie - Bonne décision	3 <sup>ème</sup> Score A : 2 Score B : 6 Choisie - Bonne décision	4 <sup>ème</sup> Score A : 1 Score B : 11 Éliminée	5 <sup>ème</sup> Score A : 3 Score B : 5 Choisie - Bonne décision
6 <sup>ème</sup> Score A : 0 Score B : 9 Éliminée	7 <sup>ème</sup> Score A : 3 Score B : 0 Éliminée	8 <sup>ème</sup> Score A : 3 Score B : 9 Choisie - Bonne décision	9 <sup>ème</sup> Score A : 1 Score B : 6 Éliminée	10 <sup>ème</sup> Score A : 3 Score B : 8 Choisie - Bonne décision

Continuer

**Essai d'entraînement n°1 / 3 - étape 1 : exploration**  
**Explorer 5 candidatures puis sélectionner les 5 meilleures parmi les 10**

Score A : 7 Score B : ? Score B à voir	Score A : 3 Score B : ? Score B à voir	Score A : 2 Score B : ? Score B caché	Score A : 10 Score B : ? Score B caché	Score A : 0 Score B : ? Score B caché
Score A : 7 Score B : ? Score B caché	Score A : 2 Score B : ? Score B à voir	Score A : 1 Score B : ? Score B à voir	Score A : 2 Score B : ? Score B à voir	Score A : 8 Score B : ? Score B caché

Valider mes demandes d'exploration et passer à l'étape de sélection

**Essai d'entraînement n°1 / 3 - étape 2 : sélection**  
**Explorer 5 candidatures puis sélectionner les 5 meilleures parmi les 10**

Score A : 7 Score B : 7 Choisie	Score A : 3 Score B : 2 Éliminée	Score A : 2 Score B : ? Éliminée	Score A : 10 Score B : ? Choisie	Score A : 0 Score B : ? Éliminée
Score A : 7 Score B : ? Choisie	Score A : 2 Score B : 7 Choisie	Score A : 1 Score B : 4 Éliminée	Score A : 2 Score B : 7 Éliminée	Score A : 8 Score B : ? Choisie

Achever le processus et enregistrer mes recrutements

**Essai d'entraînement n°1 / 3 - résultats**  
 Parmi les 5 candidatures sélectionnées, 5 sont parmi les meilleures candidatures.  
 Toutes les candidatures classées entre la 1<sup>ère</sup> et la 5<sup>ème</sup> place sont parmi les meilleures candidatures (soit 5 candidatures ou plus en cas d'égalité).  
 Les candidatures classées au-delà de la 5<sup>ème</sup> place étaient finalement les moins bonnes candidatures.

1 <sup>ère</sup> Score A : 7 Score B : 7 Choisie - Bonne décision	2 <sup>ème</sup> Score A : 3 Score B : 2 Éliminée	3 <sup>ème</sup> Score A : 2 Score B : 2 Éliminée	4 <sup>ème</sup> Score A : 10 Score B : 10 Choisie - Bonne décision	5 <sup>ème</sup> Score A : 0 Score B : 8 Éliminée
6 <sup>ème</sup> Score A : 7 Score B : 7 Choisie - Bonne décision	7 <sup>ème</sup> Score A : 2 Score B : 7 Choisie - Bonne décision	8 <sup>ème</sup> Score A : 1 Score B : 4 Éliminée	9 <sup>ème</sup> Score A : 2 Score B : 7 Éliminée	10 <sup>ème</sup> Score A : 8 Score B : 5 Choisie - Bonne décision

Continuer

## Screenshots of Education Group Intervention

### Tâche d'exploration et de sélection

**Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?**

- Page 1 -

Dans la tâche que vous venez de réaliser, l'**objectif est d'utiliser l'étape d'exploration pour augmenter vos chances de faire la meilleure sélection possible** de 5 candidatures.

Nous allons vous expliquer comment sélectionner les meilleures candidatures après avoir exploré les candidatures les plus utiles.

Ensuite, nous vous poserons quelques questions afin que vous puissiez réaliser les prochains essais en ayant une bonne compréhension de la tâche.

Vous pouvez parcourir les pages et revenir en arrière si besoin, **sauf si vous passez la 13<sup>ème</sup> et dernière page.**

Début

Page suivante (2/13)

### Tâche d'exploration et de sélection

**Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?**

- Page 2 -

Voici un exemple d'exploration...

 Score A : 3 Score B : Voir le score B	 Score A : 2 Score B : Score B masqué	 Score A : 9 Score B : Voir le score B	 Score A : 8 Score B : Voir le score B	 Score A : 5 Score B : Score B masqué
 Score A : 1 Score B : Score B masqué	 Score A : 7 Score B : Score B masqué	 Score A : 6 Score B : Voir le score B	 Score A : 0 Score B : Voir le score B	 Score A : 8 Score B : Score B masqué

Page précédente (1/13)

Page suivante (3/13)

### Tâche d'exploration et de sélection

Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?

- Page 3 -

... qui nous donne accès à 5 scores B.

 Score A : 3 Score B : 6 Choisir	 Score A : 2 Score B : ? Choisir	 Score A : 9 Score B : 2 Choisir	 Score A : 8 Score B : 5 Choisir	 Score A : 5 Score B : ? Choisir
 Score A : 1 Score B : ? Choisir	 Score A : 7 Score B : ? Choisir	 Score A : 6 Score B : 3 Choisir	 Score A : 0 Score B : 8 Choisir	 Score A : 8 Score B : ? Choisir

Page précédente (2/13)

Page suivante (4/13)

### Tâche d'exploration et de sélection

Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?

- Page 4 -

Pour la sélection, on peut remplacer les scores masqués par 5.

 Score A : 3 Score B : 6 Éliminée	 Score A : 2 Score B masqué : 5 ? Éliminée	 Score A : 9 Score B : 2 Choisie	 Score A : 8 Score B : 5 Choisie	 Score A : 5 Score B masqué : 5 ? Choisie
 Score A : 1 Score B masqué : 5 ? Éliminée	 Score A : 7 Score B masqué : 5 ? Choisie	 Score A : 6 Score B : 3 Éliminée	 Score A : 0 Score B : 8 Choisir	 Score A : 8 Score B masqué : 5 ? Choisie

Pourquoi ?

Tous les scores de 0 à 10 ont la même chance de survenir donc le score moyen est 5.

Page précédente (3/13)

Page suivante (5/13)

### Tâche d'exploration et de sélection

Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?

- Page 5 -

Concentrons-nous maintenant sur la candidature qui a le meilleur score A.

 Score A : 3 Score B : 6 Éliminée	 Score A : 2 Score B masqué : 5? Éliminée	 Score A : 9 Score B : 2 Choisie	 Score A : 8 Score B : 5 Choisie	 Score A : 5 Score B masqué : 5? Choisie
 Score A : 1 Score B masqué : 5? Éliminée	 Score A : 7 Score B masqué : 5? Choisie	 Score A : 6 Score B : 3 Éliminée	 Score A : 0 Score B : 8 Choisir	 Score A : 8 Score B masqué : 5? Choisie

On observe que son score total est de 11 et elle fait partie des 5 meilleures.

Était-ce utile d'explorer son score B ?

Non, car avec un score A de 9 et même sans connaître le score B, on était déjà presque assurés que cette candidature terminerait dans le top 5. Dans cet exemple, le score B est faible et malgré cela, cette candidature reste dans les 5 meilleures.

Obtenir l'information sur le score B n'était donc pas utile.

Page précédente (4/13)

Page suivante (6/13)

### Tâche d'exploration et de sélection

Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?

- Page 6 -

Concentrons-nous maintenant sur la candidature qui a le moins bon score A.

 Score A : 3 Score B : 6 Éliminée	 Score A : 2 Score B masqué : 5 ? Éliminée	 Score A : 9 Score B : 2 Choisie	 Score A : 8 Score B : 5 Choisie	 Score A : 5 Score B masqué : 5 ? Choisie
 Score A : 1 Score B masqué : 5 ? Éliminée	 Score A : 7 Score B masqué : 5 ? Choisie	 Score A : 6 Score B : 3 Éliminée	 Score A : 0 Score B : 8 Choisir	 Score A : 8 Score B masqué : 5 ? Choisie

On observe que son score total est de 8 et elle ne fait pas partie des 5 meilleures.

Était-ce utile d'explorer son score B ?

Non, car avec un score A de 0 et même sans connaître le score B, on était déjà presque assurés que cette candidature terminerait hors du top 5. Dans cet exemple, le score B est élevé et malgré cela, cette candidature reste hors des 5 meilleures.

Obtenir l'information sur le score B n'était donc pas utile.

En somme, il n'est pas utile d'explorer les candidatures ayant des scores A extrêmes. Le score B ne changera très probablement pas le fait que la candidature soit ou non dans le top 5.

Page précédente (5/13)

Page suivante (7/13)

### Tâche d'exploration et de sélection

**Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?**

- Page 7 -

Voici les scores A de l'exemple, classés du plus grand au plus petit.

Quels sont les scores B utiles à explorer ?

Une information est utile si la décision que l'on va prendre dépend de cette information.

Pour les candidatures qui sont pour l'instant dans le top 5, quelles sont leurs chances d'en sortir ?

Pour estimer cela, on va compter combien de **points d'avance** elles ont sur la 6<sup>ème</sup> candidature (la meilleure candidature hors top 5).

	Top 5					6 <sup>ème</sup>					
Scores A	9	8	8	7	6	5	3	2	1	0	
Points d'avance sur le 6 <sup>ème</sup>	4	3	3	2	1						

La dernière candidature du top 5 peut facilement se faire rattraper. Connaître le score B de cette candidature permet de préciser ses chances de se faire dépasser.

Page précédente (6/13)

Page suivante (8/13)

### Tâche d'exploration et de sélection

**Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?**

- Page 8 -

Voici les scores A de l'exemple, classés du plus grand au plus petit.

Quels sont les scores B utiles à explorer ?

Une information est utile si la décision que l'on va prendre dépend de cette information.

Pour les candidatures qui sont pour l'instant hors top 5, quelles sont leurs chances d'y entrer ?

Pour estimer cela, on va compter combien de **points de retard** elles ont sur le 5<sup>ème</sup> candidature (la dernière candidature du top 5).

					5 <sup>ème</sup> ▼	<b>Hors Top 5</b>				
Scores A	9	8	8	7	6	5	3	2	1	0
Points de retard sur le 5 <sup>ème</sup>						1	3	4	5	6

La première candidature hors top 5 peut facilement rattraper son retard. Connaître le score B de cette candidature permet de préciser ses chances de rattraper le top 5.

En somme, il est utile d'explorer les candidatures qui ont de fortes chances de basculer dans le top 5 ou en dehors du top 5, étant donné leur score A. Le score B peut changer le fait qu'une candidature soit ou non dans le top 5 final.

Page précédente (7/13)

Page suivante (9/13)

### Tâche d'exploration et de sélection

Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?

- Page 9 -

En prenant les candidatures qui ont le moins de points d'avance ou de retard, on veut donc explorer les 5 candidatures qui auraient le plus de chance de basculer. Dans l'exemple ci-dessous, cela inclut les 3 candidatures en bleu (scores A de 7, 6 et 5), et il restera ensuite à choisir 2 candidatures parmi les 3 candidatures en violet (scores A de 8, 8 et 3). Comment faire ce choix ?

	Top 5					Hors Top 5				
Scores A	9	8	8	7	6	5	3	2	1	0
Points d'avance sur le 6 <sup>ème</sup>	4	3	3	2	1					
Points de retard sur le 5 <sup>ème</sup>						1	3	4	5	6

Comment départager les égalités sur les points de retard / d'avance ?

On utilise cette fois les rangs. Les scores et les rangs sont liés mais ils sont assez différents pour départager la plupart des égalités.

Page précédente (8/13)

Page suivante (10/13)

### Tâche d'exploration et de sélection

Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?

- Page 10 -

Sur la base des points d'avance et de retard, on veut donc explorer les candidatures ayant les scores A de 7, 6 et 5. Il nous reste à choisir deux candidatures à explorer parmi les deux candidatures ayant des scores A de 8 et celle ayant un 3.

Pour les candidatures à départager :

- Si elles sont **dans le top 5**, on regarde le nombre de **rangs d'avance** par rapport au **meilleur rang hors top 5** (ici, le 6<sup>ème</sup> rang).
- Si elles sont **hors du top 5**, on regarde le nombre de **rangs de retard** par rapport au **moins bon rang du top 5** (ici, le 5<sup>ème</sup> rang).

	Top 5					Hors Top 5				
	1 <sup>er</sup>	2 <sup>èmes</sup>	4 <sup>ème</sup>	5 <sup>ème</sup>	6 <sup>ème</sup>	7 <sup>ème</sup>	8 <sup>ème</sup>	9 <sup>ème</sup>	10 <sup>ème</sup>	
Scores A	9	8	8	7	6	5	3	2	1	0
Rangs d'avance sur le 6 <sup>ème</sup>		4								
Rangs de retard sur le 5 <sup>ème</sup>						2				
Points d'avance sur le 6 <sup>ème</sup>	4	3	3	2	1					
Points de retard sur le 5 <sup>ème</sup>						1	3	4	5	6

Sur la base des rangs, on privilégie donc d'explorer la candidature ayant un score A de 3.

Page précédente (9/13)

Page suivante (11/13)

### Tâche d'exploration et de sélection

Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?

- Page 11 -

Sur la base des points d'avance et de retard, on veut explorer les candidatures ayant les scores A de 7, 6 et 5. Sur la base des rangs, on privilégie d'explorer la candidature ayant un score A de 3.

Il nous reste à choisir une candidature à explorer parmi les deux candidatures ayant des scores A de 8.

Pour toutes les égalités restantes, on choisit au hasard car rien d'autre ne permet de différencier les candidatures ! On choisit donc au hasard une des deux candidatures ayant un score A de 8 pour compléter notre exploration.

	Top 5					Hors Top 5				
	1 <sup>er</sup>	2 <sup>èmes</sup>	4 <sup>ème</sup>	5 <sup>ème</sup>	6 <sup>ème</sup>	7 <sup>ème</sup>	8 <sup>ème</sup>	9 <sup>ème</sup>	10 <sup>ème</sup>	
Scores A	9	8	8	7	6	5	3	2	1	0
Rangs d'avance sur le 6 <sup>ème</sup>		4								
Rangs de retard sur le 5 <sup>ème</sup>						2				
Points d'avance sur le 6 <sup>ème</sup>	4	3	3	2	1					
Points de retard sur le 5 <sup>ème</sup>						1	3	4	5	6

Page précédente (10/13)

Page suivante (12/13)

### Tâche d'exploration et de sélection

**Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?**

- Page 12 -

Les candidatures à explorer dans cet exemple sont donc celles ayant les scores A de 7, 6, 5, 3 et une des deux ayant un 8 (choisie au hasard).

#### Exploration

Scores A	9	8	8	7	6	5	3	2	1	0
----------	---	---	---	---	---	---	---	---	---	---

Concrètement, voici l'exploration :

 Score A : 3 Score B : <a href="#">Voir le score B</a>	 Score A : 2 Score B : <a href="#">Score B masqué</a>	 Score A : 9 Score B : <a href="#">Score B masqué</a>	 Score A : 8 Score B : <a href="#">Voir le score B</a>	 Score A : 5 Score B : <a href="#">Voir le score B</a>
 Score A : 1 Score B : <a href="#">Score B masqué</a>	 Score A : 7 Score B : <a href="#">Voir le score B</a>	 Score A : 6 Score B : <a href="#">Voir le score B</a>	 Score A : 0 Score B : <a href="#">Score B masqué</a>	 Score A : 8 Score B : <a href="#">Score B masqué</a>

[Page précédente \(11/13\)](#)

[Page suivante \(Dernière\)](#)

## Tâche d'exploration et de sélection

**Quelles sont les 5 options à explorer pour trouver les 5 meilleures candidatures ?**

- Page 13 -

*Récapitulatif :*

### 1. Exploration :

Explorer les candidatures qui ont le plus de chances de rentrer ou de sortir dans le top 5 selon leur score A (et leur rang provisoire).

*Ce sont les candidatures pour lesquelles les scores B seront les plus utiles, car elles peuvent encore basculer dans le top 5 ou en être éjectées.*

### 2. Sélection :

Remplacer les scores B masqués par 5.

Choisir les candidatures qui ont les meilleurs scores totaux (score A + score B).

*Ce sont les candidatures qui ont le plus de chances d'être réellement parmi les 5 meilleures, étant donné les scores A et B à votre disposition.*

**Attention : en passant cette page, nous ne pourrez plus revenir en arrière.**

Page précédente (12/13)

Page suivante (Suite)

## Entraînement

Vous avez **2 essais d'entraînement guidés** afin de vous familiariser avec les explications sur l'exploration et la sélection. Votre objectif est d'appliquer au mieux les conseils que vous venez de recevoir. **Vos réponses ne seront pas comptabilisées.**

Appuyez sur « Commencer » pour lancer l'entraînement.

Commencer

## Tâche d'exploration et de sélection

### Question n°1 / 5

S'il vous plaît, veuillez corriger les erreurs.

Avant de commencer la partie 2, merci de répondre aux affirmations suivantes par Vrai ou Faux. Cliquez sur "Question suivante" pour valider votre réponse.

En cas de mauvaise réponse, vous verrez apparaître un rappel de ce qui vous a sûrement échappé. Il faudra corriger votre réponse avant de pouvoir passer à la page suivante. Vous pouvez consulter le "Rappel" si besoin.

« Explorer le score B de la candidature dont le score A est classé 1<sup>er</sup> est souvent très utile pour trouver les 5 meilleures candidatures lors de la sélection. »

Vrai  Faux

Une candidature qui a le meilleur score est presque certaine d'être dans les 5 meilleures du classement final. Quel que soit son score B, il ne vous sera pas très utile pour trouver les 5 meilleures candidatures.

#### ▼ Rappel

##### 1. Exploration :

Explorer les candidatures qui ont le plus de chances de rentrer ou de sortir dans le top 5 selon leur score A (et leur rang provisoire.)

*Ce sont les candidatures pour lesquelles les scores B seront les plus utiles, car elles peuvent encore basculer dans le top 5 ou en être éjectées.*

##### 2. Sélection :

Remplacer les scores B masqués par 5.

Choisir les candidatures qui ont les meilleurs scores totaux (score A + score B).

*Ce sont les candidatures qui ont le plus de chances d'être réellement parmi les 5 meilleures, étant donné les scores A et B à votre disposition.*

Question suivante

**Question n°2 / 5**

« Explorer le score B d'une candidature qui a un score A classé 8<sup>ème</sup> et qui est très éloigné du score A de la candidature classée 6<sup>ème</sup> n'est pas vraiment utile pour trouver les 5 meilleures candidatures lors de la sélection. »

Vrai  Faux

Une candidature qui a un classement faible et un score éloigné des scores en milieu de classement est presque certaine de ne pas être dans les 5 meilleures du classement final. Quel que soit son score B, il ne vous sera pas très utile pour trouver les 5 meilleures candidatures.

**Question n°3 / 5**

« Explorer le score B d'une candidature dont le score A est au milieu du classement provisoire est très souvent utile pour trouver les 5 meilleures candidatures lors de la sélection. »

Vrai  Faux

Il est toujours utile d'explorer une candidature dont le score A a un classement moyen car c'est une candidature pour laquelle le score B sera très utile pour augmenter vos chances de trouver quelles candidatures sont parmi les 5 meilleures.

**Question n°4 / 5**

« Une candidature ayant les scores 6 et 7 aura forcément un score total supérieur à une candidature ayant un score A de 5 et un score B masqué. »

Vrai  Faux

Un score total de 13 (6+7) est inférieur aux scores de 15 et 14 qui sont possibles si le un score B masqué est 10 ou 9. Il n'y a aucune garantie que ces scores surviennent, mais ils sont possibles donc une candidature ayant des scores 6 et 7, n'est pas assurée à 100% d'avoir un meilleur score total qu'une candidature ayant un score A de 5 et un score B masqué.

**Question n°5 / 5**

« Un score B masqué peut être remplacé par un 5 afin de faire la meilleure sélection possible, mais il n'y a aucune garantie que le score B soit réellement un 5. »

Vrai  Faux

Un score B masqué peut être n'importe quel nombre entier entre 0 et 10. Tous les scores entre 0 et 10 ont la même probabilité de survenir. En conséquence, la valeur moyenne d'un score B masqué peut être estimée à 5. Ce n'est qu'une estimation qui permet de faire une sélection mais pas nécessairement la valeur réelle de chaque score B masqué.

Essai d'entrainement n°1 / 2

Explorer 5 candidatures puis sélectionner les 5 meilleures parmi les 10

 Score A : <b>10</b> Score B :	 Score A : <b>9</b> Score B :	 Score A : <b>8</b> Score B :	 Score A : <b>7</b> Score B :	 Score A : <b>6</b> Score B :
 Score A : <b>5</b> Score B :	 Score A : <b>4</b> Score B :	 Score A : <b>1</b> Score B :	 Score A : <b>0</b> Score B :	 Score A : <b>0</b> Score B :

Quel est le 5<sup>ème</sup> meilleur score A ? (Indiquez uniquement un nombre)

Quel est le 6<sup>ème</sup> meilleur score A ? (Indiquez uniquement un nombre)

Commencer

Essai d'entraînement n°2 / 2

Explorer 5 candidatures puis sélectionner les 5 meilleures parmi les 10

 Score A : <b>1</b> Score B :	 Score A : <b>7</b> Score B :	 Score A : <b>7</b> Score B :	 Score A : <b>9</b> Score B :	 Score A : <b>0</b> Score B :
 Score A : <b>9</b> Score B :	 Score A : <b>5</b> Score B :	 Score A : <b>7</b> Score B :	 Score A : <b>2</b> Score B :	 Score A : <b>3</b> Score B :

Quel est le dernier rang inclus dans le top 5 ? (Indiquez uniquement un nombre)

Quel est le premier rang en dehors du top 5 ? (Indiquez uniquement un nombre)

Commencer

## Entraînement

Fin de l'entraînement

Bravo ! Vous avez réussi à compléter l'entraînement.

Voici quelques questions avant de commencer la partie 2.

Passer aux questions

### Essai d'entrainement n°1 / 2 - étape 1 : exploration

Explorer 5 candidatures puis sélectionner les 5 meilleures parmi les 10

 Score A : <b>10</b> Score B : Score B à voir	 Score A : <b>9</b> Score B : Score B à voir	 Score A : <b>8</b> Score B : Score B caché	 Score A : <b>7</b> Score B : Score B à voir	 Score A : <b>6</b> Score B : Score B caché
 Score A : <b>5</b> Score B : Score B caché	 Score A : <b>4</b> Score B : Score B à voir	 Score A : <b>1</b> Score B : Score B à voir	 Score A : <b>0</b> Score B : Score B caché	 Score A : <b>0</b> Score B : Score B caché

#### ▼ Commentaires sur vos demandes d'exploration (5)

**Score A de 10** : Très éloigné du 6<sup>ème</sup> meilleur score A. Certaines candidatures sont plus intéressantes à explorer.

**Score A de 9** : Éloigné du score 6<sup>ème</sup> meilleur score A. Certaines candidatures sont plus intéressantes à explorer.

**Score A de 7** : Avec une faible distance du 6<sup>ème</sup> meilleur score A, le score B sera très utile pour prendre une décision.

**Score A de 4** : Avec une faible distance du 5<sup>ème</sup> meilleur score A, le score B sera très utile pour prendre une décision.

**Score A de 1** : Très éloigné du 5<sup>ème</sup> meilleur score A. Certaines candidatures sont plus intéressantes à explorer.

**Vous devez modifier votre exploration (consultez les commentaires)**

### Essai d'entrainement n°1 / 2 - étape 1 : exploration

Explorer 5 candidatures puis sélectionner les 5 meilleures parmi les 10

 Score A : <b>10</b> Score B : Score B caché	 Score A : <b>9</b> Score B : Score B caché	 Score A : <b>8</b> Score B : Score B à voir	 Score A : <b>7</b> Score B : Score B à voir	 Score A : <b>6</b> Score B : Score B à voir
 Score A : <b>5</b> Score B : Score B à voir	 Score A : <b>4</b> Score B : Score B à voir	 Score A : <b>1</b> Score B : Score B caché	 Score A : <b>0</b> Score B : Score B caché	 Score A : <b>0</b> Score B : Score B caché

#### ▼ Commentaires sur vos demandes d'exploration (5)

**Score A de 8** : Assez proche du 6<sup>ème</sup> meilleur score A. Le score B sera utile pour prendre une décision.

**Score A de 7** : Avec une faible distance du 6<sup>ème</sup> meilleur score A, le score B sera très utile pour prendre une décision.

**Score A de 6** : Avec une très faible distance du 6<sup>ème</sup> meilleur score A, le score B sera très utile pour prendre une décision.

**Score A de 5** : Avec une très faible distance du 5<sup>ème</sup> meilleur score A, le score B sera très utile pour prendre une décision.

**Score A de 4** : Avec une faible distance du 5<sup>ème</sup> meilleur score A, le score B sera très utile pour prendre une décision.

Valider mes demandes d'exploration et passer à l'étape de sélection

**Essai d'entrainement n°1 / 2 - étape 2 : sélection**

**Explorer 5 candidatures puis sélectionner les 5 meilleures parmi les 10**

 Score A : <b>10</b> Score B : ? Choisie	 Score A : <b>9</b> Score B : ? Choisie	 Score A : <b>8</b> Score B : <b>4</b> Choisie	 Score A : <b>7</b> Score B : <b>7</b> Choisie	 Score A : <b>6</b> Score B : <b>0</b> Éliminée
 Score A : <b>5</b> Score B : <b>5</b> Choisie	 Score A : <b>4</b> Score B : <b>9</b> Éliminée	 Score A : <b>1</b> Score B : ? Éliminée	 Score A : <b>0</b> Score B : ? Éliminée	 Score A : <b>0</b> Score B : ? Éliminée

▼ **Commentaires sur vos choix de recrutement (5)**

**Scores de 10 et masqué :** Avec un score B moyen de 5, cette candidature aurait un score total de 15. Ce serait la 1<sup>ère</sup> candidature du classement final et donc une bonne décision de recrutement.

**Scores de 9 et masqué :** Avec un score B moyen de 5, cette candidature aurait un score total de 14. Ce serait la 2<sup>ème</sup> candidature du classement final et donc une bonne décision de recrutement.

**Scores de 8 et 4 :** Le score total de 12 devrait permettre à cette candidature d'espérer se classer 5<sup>ème</sup> et serait donc une bonne décision de recrutement.

**Scores de 7 et 7 :** Le score total de 14 devrait permettre à cette candidature d'espérer se classer 2<sup>ème</sup> et serait donc une bonne décision de recrutement.

**Scores de 5 et 5 :** Le score total de 10 ne devrait pas permettre à cette candidature d'être dans le top 5 et serait donc une mauvaise décision de recrutement.

**Vous devez modifier votre sélection (consultez les commentaires)**

**Essai d'entrainement n°1 / 2 - étape 2 : sélection**

**Explorer 5 candidatures puis sélectionner les 5 meilleures parmi les 10**

 Score A : <b>10</b> Score B : ? Choisie	 Score A : <b>9</b> Score B : ? Choisie	 Score A : <b>8</b> Score B : <b>4</b> Choisie	 Score A : <b>7</b> Score B : <b>7</b> Choisie	 Score A : <b>6</b> Score B : <b>0</b> Éliminée
 Score A : <b>5</b> Score B : <b>5</b> Éliminée	 Score A : <b>4</b> Score B : <b>9</b> Choisie	 Score A : <b>1</b> Score B : ? Éliminée	 Score A : <b>0</b> Score B : ? Éliminée	 Score A : <b>0</b> Score B : ? Éliminée

▼ **Commentaires sur vos choix de recrutement (5)**

**Scores de 10 et masqué :** Avec un score B moyen de 5, cette candidature aurait un score total de 15. Ce serait la 1<sup>ère</sup> candidature du classement final et donc une bonne décision de recrutement.

**Scores de 9 et masqué :** Avec un score B moyen de 5, cette candidature aurait un score total de 14. Ce serait la 2<sup>ème</sup> candidature du classement final et donc une bonne décision de recrutement.

**Scores de 8 et 4 :** Le score total de 12 devrait permettre à cette candidature d'espérer se classer 5<sup>ème</sup> et serait donc une bonne décision de recrutement.

**Scores de 7 et 7 :** Le score total de 14 devrait permettre à cette candidature d'espérer se classer 2<sup>ème</sup> et serait donc une bonne décision de recrutement.

**Scores de 4 et 9 :** Le score total de 13 devrait permettre à cette candidature d'espérer se classer 4<sup>ème</sup> et serait donc une bonne décision de recrutement.

Achever le processus et enregistrer mes recrutements

## Screenshots of Control Group Intervention

### Tâche d'exploration et de sélection

#### Rappel des instructions

- Page 1 -

Vous venez de compléter les 10 premiers essais de l'expérience.

Avant de poursuivre, nous vous invitons à relire **attentivement** les instructions.

Après cette relecture, nous vous poserons quelques questions afin que vous puissiez réaliser les prochains essais en ayant une bonne compréhension de la tâche.

Vous pouvez parcourir les pages et revenir en arrière si besoin, **sauf si vous passez la 8<sup>ème</sup> et dernière page.**

Début

Page suivante (2/8)

### Tâche d'exploration et de sélection

#### Rappel des instructions

- Page 2 -

Vous allez prendre part à une tâche dont l'objectif est de réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Lors de chaque processus, votre mission est de **sélectionner les 5 meilleures candidatures parmi une liste de 10 candidatures.**

Pour chaque candidature, vous aurez accès à **un score A allant de 0 (inclus) à 10 (inclus)** correspondant à une évaluation d'une compétence A de chaque personne.

Parmi les 10 candidatures, vous pourrez alors indiquer **5 candidatures** pour lesquels vous aurez accès à **un score B** correspondant à une évaluation d'une compétence B de chaque personne.

Enfin, vous devrez prendre une décision de recrutement, à savoir **choisir selon vous les 5 meilleures candidatures** parmi les 10 candidatures.

Page précédente (1/8)

Page suivante (3/8)

## Tâche d'exploration et de sélection

### Rappel des instructions

- Page 3 -

Lors de la première étape (l'exploration), vous verrez apparaître simultanément à l'écran les 10 candidatures ainsi qu'un score A allant de 0 à 10, correspondant à une évaluation de la compétence A. Plus le score A est élevé, plus la personne a été évaluée comme possédant la compétence A, nécessaire pour le poste en question.

Les candidatures seront présentés de la manière suivante :

 Score A : <b>5</b> Score B : <input type="button" value="Voir score B"/>	 Score A : <b>7</b> Score B : <input type="button" value="Voir score B"/>	 Score A : <b>3</b> Score B : <input type="button" value="Voir score B"/>	 Score A : <b>3</b> Score B : <input type="button" value="Voir score B"/>	 Score A : <b>3</b> Score B : <input type="button" value="Voir score B"/>
 Score A : <b>4</b> Score B : <input type="button" value="Voir score B"/>	 Score A : <b>10</b> Score B : <input type="button" value="Voir score B"/>	 Score A : <b>7</b> Score B : <input type="button" value="Voir score B"/>	 Score A : <b>9</b> Score B : <input type="button" value="Voir score B"/>	 Score A : <b>9</b> Score B : <input type="button" value="Voir score B"/>

Pour chaque candidature, vous pouvez cliquer sur le bouton "**Voir score B**" si vous souhaitez l'**explorer**, c'est-à-dire **voir le score d'évaluation de la compétence B de la personne correspondante** lors de l'étape suivante.

Vous devez indiquer **exactement 5 candidatures à explorer** : vous ne pourrez passer à l'étape suivante que lorsque vous aurez indiqué 5 personnes dont vous voulez voir le score B.

Pour chaque candidature, **vous pouvez changer d'avis** autant de fois que vous le souhaitez en cliquant de nouveau sur le même bouton.

Si vous avez déjà indiqué 5 candidatures à explorer, vous devez en retirer une avant de pouvoir en ajouter une autre.

Le nombre de clics que vous effectuez ainsi que vos temps de réponse sont enregistrés, mais cela n'a pas d'influence sur vos gains potentiels.

Vos demandes d'exploration sont enregistrées et définitives lorsque vous cliquez sur le bouton "Valider mes demandes d'exploration et passer à l'étape de sélection", en bas de page.

Page précédente (2/8)

Page suivante (4/8)

### Tâche d'exploration et de sélection

#### Rappel des instructions

- Page 4 -

Lors de la deuxième étape (la sélection), vous verrez les 10 mêmes candidatures avec le score A que vous aviez déjà vu précédemment et le **score B pour les 5 candidatures que vous avez demandé à explorer lors de l'étape précédente**. Pour les autres personnes, le score B est remplacé par un " ? " pour indiquer qu'il **existe aussi**, mais que vous n'en avez pas connaissance. Plus le score B est élevé, plus la personne a été évaluée comme possédant la compétence B, nécessaire pour le poste en question.

Les candidatures seront présentées de la manière suivante :

 Score A : <b>5</b> Score B : <b>7</b> Choisir	 Score A : <b>7</b> Score B : ? Choisir	 Score A : <b>3</b> Score B : <b>3</b> Choisir	 Score A : <b>3</b> Score B : ? Choisir	 Score A : <b>3</b> Score B : <b>3</b> Choisir
 Score A : <b>4</b> Score B : ? Choisir	 Score A : <b>10</b> Score B : <b>1</b> Choisir	 Score A : <b>7</b> Score B : ? Choisir	 Score A : <b>9</b> Score B : <b>2</b> Choisir	 Score A : <b>9</b> Score B : ? Choisir

L'étape fonctionne de la même manière que l'étape précédente.

Pour chaque candidature, vous pouvez cliquer sur le bouton "Choisir" si vous souhaitez recruter cette personne.

**Votre objectif est de sélectionner les 5 candidatures qui ont les meilleurs scores en faisant l'addition de leurs deux scores (A et B). Pour déterminer les meilleures candidatures, les deux scores de chaque personne sont pris en compte, que vous ayez décidé d'explorer le score B ou non.**

Vous devez sélectionner **exactement 5 candidatures** : vous ne pourrez valider votre recrutement que lorsque vous aurez recruté 5 personnes.

**Toutes les candidatures peuvent être choisies** : vous pouvez choisir les candidatures pour lesquelles vous avez vu un seul score comme les candidatures pour lesquelles vous avez vu les deux scores.

Pour chaque candidature, **vous pouvez changer d'avis** autant de fois que vous le souhaitez en cliquant de nouveau sur le même bouton.

Si vous avez déjà indiqué 5 candidatures à sélectionner, vous devez en retirer une avant de pouvoir en choisir une autre.

Le nombre de clics que vous effectuez ainsi que vos temps de réponse sont enregistrés, mais cela n'a pas d'influence sur vos gains potentiels.

Vos décisions de recrutement sont enregistrées et définitives lorsque vous cliquez sur le bouton "Achever le processus et enregistrer mes recrutements", en bas de page.

Page précédente (3/8)

Page suivante (5/8)

### Tâche d'exploration et de sélection

#### Rappel des instructions

- Page 5 -

À la suite de votre recrutement, vous aurez une indication sur le nombre de candidatures choisies qui étaient effectivement parmi les meilleures candidatures.

Le résultat sera présenté sous la forme suivante :

Parmi les 5 candidatures sélectionnées, 4 étaient parmi les meilleures candidatures.

Pour vous permettre de mieux comprendre comment est calculé le résultat, voici une vue plus complète à laquelle vous aurez accès uniquement pendant l'entraînement :

Parmi les 5 candidatures sélectionnées, 4 étaient parmi les meilleures candidatures.

Les candidatures sont classées dans l'ordre décroissant selon la somme de leurs deux scores. La candidature ayant le score total le plus élevé est première et la candidature ayant le score total le plus faible est dernière.

Toutes les candidatures classées entre la 1<sup>ère</sup> et la 5<sup>ème</sup> place sont parmi les meilleures candidatures (soit 5 candidatures ou plus en cas d'égalité).

Les candidatures classées au-delà de la 5<sup>ème</sup> place étaient les moins bonnes candidatures, c'est-à-dire celles qu'il fallait éliminer.

 3 <sup>ème</sup> Score A : 5 Score B : 7 Choisie - Bonne décision	 6 <sup>ème</sup> Score A : 7 Score B : 3 Choisie - Mauvaise décision	 8 <sup>ème</sup> Score A : 3 Score B : 3 Éliminée	 6 <sup>ème</sup> Score A : 3 Score B : 7 Éliminée	 8 <sup>ème</sup> Score A : 3 Score B : 3 Éliminée
 10 <sup>ème</sup> Score A : 4 Score B : 0 Éliminée	 4 <sup>ème</sup> Score A : 10 Score B : 1 Choisie - Bonne décision	 1 <sup>ère</sup> Score A : 7 Score B : 8 Éliminée	 4 <sup>ème</sup> Score A : 9 Score B : 2 Choisie - Bonne décision	 2 <sup>ème</sup> Score A : 9 Score B : 5 Choisie - Bonne décision

Parmi les 5 candidatures sélectionnées, choisir la deuxième candidature de la rangée du haut était une mauvaise décision, car elle est classée à la 6<sup>ème</sup> place, c'est-à-dire après la 5<sup>ème</sup> place.

À l'inverse, la candidature située au milieu de la rangée du bas était 1<sup>ère</sup>, mais n'a pas été choisie.

En cas d'égalité, les candidatures sont considérées ex-æquo et sont classées à la même place. Si elles sont entre la 1<sup>ère</sup> et la 5<sup>ème</sup> place, elles sont considérées comme des bonnes réponses indifféremment.

Pendant l'expérience, le résultat de vos décisions de recrutement est affiché pendant 5 secondes avant que le processus de recrutement suivant ne commence.

Page précédente (4/8)

Page suivante (6/8)

## Tâche d'exploration et de sélection

### Rappel des instructions

- Page 6 -

Les scores A et B des candidatures sont des **nombre entiers tirés aléatoirement** de 0 à 10. Ils ont tous la **même chance** d'être tirés au sort.

Ainsi, chaque score est **indépendant** et a la **même probabilité de survenir** quelque soit les autres scores. En d'autres termes, tous les scores A sont indépendants, tous les scores B sont indépendants, et les scores A et B de chaque candidature sont indépendants.

Page précédente (5/8)

Page suivante (7/8)

## Tâche d'exploration et de sélection

### Rappel des instructions

- Page 7 -

*Récapitulatif :*

Vous allez réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Chaque processus de recrutement considère **10 candidatures** pour le poste.

Pour chaque processus, **votre objectif est de sélectionner les 5 meilleures candidatures**.

1) Étape d'exploration : Après avoir consulté un score d'évaluation d'une compétence A pour toutes les candidatures, vous devez **indiquer 5 candidatures pour lesquelles vous souhaitez voir un score d'évaluation d'une compétence B**.

Chaque score est compris entre **0 (inclus) et 10 (inclus)** et plus il est élevé, plus la personne est évaluée comme étant **compétente**.

2) Étape de sélection : Vous devez alors **recruter les 5 personnes qui sont au total les plus compétentes (score A + score B)**.

**Toutes les candidatures peuvent être choisies** : vous pouvez choisir les candidatures pour lesquelles vous avez vu un seul score comme les candidatures pour lesquelles vous avez vu les deux scores.

Tous les scores A sont indépendants, tous les scores B sont indépendants, et les scores A et B pour chaque candidature sont indépendants.

Vous obtiendrez des gains pour la réalisation de la tâche et en fonction de la qualité **d'un de vos recrutements tiré au hasard**. Chaque candidature choisie qui est effectivement parmi les meilleures candidatures vous rapporte **0,50€**, soit un gain total maximum de **2,50€**.

Page précédente (6/8)

Page suivante (Dernière)

## Tâche d'exploration et de sélection

### Rappel des instructions

- Page 8 -

Vous avez à nouveau 3 essais d'entraînement durant lesquelles vos réponses ne sont pas comptabilisées. Les questions concernant les instructions seront présentées juste après les 3 essais d'entraînement.

**Attention : en passant cette page, nous ne pourrez plus revenir en arrière.**

Page précédente (7/8)

Entraînement

## Instructions

### Question n°1 / 5

S'il vous plaît, veuillez corriger les erreurs.

Avant de commencer la partie suivante, merci de répondre aux affirmations suivantes concernant les instructions de la tâche par Vrai ou Faux. Cliquez sur "Question suivante" pour valider votre réponse.

En cas de mauvaise réponse, vous verrez apparaître un rappel de l'instruction qui vous a sûrement échappé. Il faudra corriger votre réponse avant de pouvoir passer à la page suivante. Vous pouvez consulter le "Rappel des instructions" si besoin.

« Il est interdit de sélectionner une candidature qui n'a pas été explorée. »

Vrai  Faux

Il est autorisé de sélectionner toutes les candidatures, y compris une candidature qui n'a pas été explorée.

#### ▼ Rappel des instructions

Vous allez réaliser **10 processus de recrutement**, indépendants les uns des autres, pour divers postes relatifs à des métiers (qui ne vous sont pas précisés).

Chaque processus de recrutement considère **10 candidatures** pour le poste.

Pour chaque processus, **votre objectif est de sélectionner les 5 meilleures candidatures.**

1) Étape d'exploration : Après avoir consulté un score d'évaluation d'une compétence A pour toutes les candidatures, vous devez **indiquer 5 candidatures pour lesquelles vous souhaitez voir un score d'évaluation d'une compétence B.**

Chaque score est compris entre **0 (inclus) et 10 (inclus)** et plus il est élevé, **plus** la personne est évaluée comme étant **compétente.**

2) Étape de sélection : Vous devez alors **recruter les 5 personnes qui sont au total les plus compétentes (score A + score B).**

Tous les scores A sont indépendants, tous les scores B sont indépendants, et les scores A et B pour chaque candidature sont indépendants.

Question suivante

### Question n°2 / 5

« Lors de la sélection, il est autorisé de recruter une candidature ayant un score total inférieur à 10. »

Vrai  Faux

Toutes les candidatures peuvent être recrutées. L'idéal étant que vous pensiez qu'elles ont une chance de faire partie des meilleures

### Question n°3 / 5

« Lorsque le score B d'une candidature n'est pas exploré, il peut être différent du score A de la candidature. »

Vrai  Faux

Le score B d'une candidature est un nombre entier entre 0 et 10. Il y a donc 1 chance sur 11 que le score B soit identique au score A, mais le score A n'a pas d'influence sur le score B, qu'il soit exploré ou non.

### Question n°4 / 5

« Il est possible d'explorer toutes les candidatures plutôt que le nombre de candidatures indiqué dans les consignes si toutes les candidatures ont le même score. »

Vrai  Faux

Quelque soit les scores présentés, vous devez explorer le nombre de candidatures indiqué dans les instructions.

### Question n°5 / 5

« Une candidature qui a un score A de 6 et un score B de 5 est meilleure qu'une candidature qui a un score A de 8 et un score B de 4. »

Vrai  Faux

Une candidature qui a les scores 6 et 5 a un score total de 11, alors qu'une candidature qui a les scores 8 et 4 a un score total de 12.

### Advice given by participants to help other future participants

At the end of the task, participants were invited to provide advice on completing the task. The instruction was: "What advice would you give to help others who have to complete the same task?"

The table below summarises all participants who provided a response. Performance is calculated as the average of their performance per block divided by the average performance of the optimal strategy per block. The quality of their exploration is measured as the average alignment of their exploration with that of the optimal exploration across the blocks. The group indicates whether participants were in the education condition (explanation of the optimal strategy provided after block 1) or the control condition (no explanation of the optimal strategy).

Participant	Advice	Perf	Explo_Quali	Group
089di1uf	vérifiez vos résultats, on peut oublier des chiffres	97.59091	84.66667	Education
bdbit5uu	LIRE LES EXPLICATIONS	93.42378	51.66667	Education
fyh5afqh	Réfléchissez bien	93.70896	56.50000	Education
my3pr692	Prendre le temps de bien analyser les données	101.37463	66.33333	Education
z7ej5q8x	Lire attentivement les instructions et explications. Ne pas sélectionner les candidatures par hasard et utiliser les explications en partie 1 pour prendre toutes les décisions dans chaque partie. Enfin, adapter la méthode expliquée en partie 1 aux conditions des parties 3 et 4.	102.65517	84.00000	Education
i632d5cb	Rester concentrer sur les chiffres et les instructions	90.38208	45.33333	Education
rp8oue3m	Se faire confiance	97.18032	72.50000	Education
w4x264m6	réfléchir	100.23202	62.83333	Education
3qs8ledz	Prendre son temps si l on en a besoin.	98.26599	75.66667	Education
hjrnefot	Bien comprendre la partie sur les "top" car cela aide beaucoup (après la 1ère partie de l expérience)	97.02381	70.00000	Education
uhzt9rcj	ne pas voir le score B des gros chiffres au début, plutôt voir ceux du milieu	95.15763	64.66667	Education
6lx99ov7	peut être un graphique plus adapté au début des choix	96.55742	56.33333	Education
6v4npj76	concentration et un endroit calme	95.96327	69.50000	Education
hvtptbzz	bon courage	93.19466	52.33333	Education
oy00h9zs	de prendre son temps et si nécessaire lister à côté l ordre des valeurs de A pour mieux avoir l information	100.00000	81.50000	Education
qlv108mr	bien lire les consignes	93.50630	62.16667	Education
vhb2gjq8	Aucun, très clair	98.83721	71.16667	Education
8xavolz	Une solution moins optimale, mais plus simple, est d'explorer les candidatures avec les scores A les plus élevés et ensuite de comparer les scores A + B explorés avec les scores A + 5 comme espérance pour les scores B que l'on ne connaît pas. C'est moins optimal, mais cela évite les erreurs de jugements pour définir quelles sont les candidatures ayant le plus de chances de rentrer ou de sortir du top à retenir et plus rapide.	95.48392	69.66667	Education
dvdab6w6	non - la chance peut jouer! ou l instinct!	91.35375	56.83333	Education
q64kml7n	réfléchissez, lisez les consignes	100.94104	79.33333	Education
16ylc5sz	Bien lire les consignes et l entraînement	90.72987	29.83333	Control

<b>Participant</b>	<b>Advice</b>	<b>Perf</b>	<b>Explo_Quali</b>	<b>Group</b>
1h9y2say	Vraiment le faire dans un endroit calme, et pouvoir calculer tranquillement	97.02267	73.50000	Control
7h9vndks	d être concentrer sur le changement de parties.	94.36191	40.00000	Control
cf92wz2n	concentration et probas	96.36931	31.83333	Control
e772qlj6	choisir ce qui ont des probabilités A moyennes, ni trop petite ni trop haut	87.60725	50.33333	Control
fkycr40d	oui de bien réfléchir avant de regarder les scores	91.40691	48.16667	Control
hrj4gddj	Pas vraiment. Juste, faites-vous confiance !	93.15406	50.50000	Control
o1v5kmag	Bien choisir les profiles que l on regarde	85.03867	51.16667	Control
tx9kd8pw	Bien comprendre les notions de probabilités en amont	90.41411	37.00000	Control
w9wdqmeq	Réfléchissez bien à la balance nombre de découvertes/ choix pour trouver quelles explorations sont les plus pertinentes	95.25802	63.66667	Control
gew7v14k	Essayer d explorer au max les personnes ayant des scores médiocres et très proches de l un de l autre au lieu de gaspiller l exploration sur les personnes ayant des scores très surs.	93.58360	54.83333	Control
l3thwu3k	Bien évaluer des candidatures dont les partie A présente des valeurs "moyennes" du point de vue de l échantillon global.	94.78426	75.16667	Control
rr3970nd	de rester concentré et de faire cette expérience au calme	96.32253	64.66667	Control
t6tqazwl	Si augmentation de trouver avec réduction d explorer alors prés sélectionner les valeurs proche de 10 donc ne pas les explorer	92.34668	32.83333	Control
1c011qms	IL FAUT PAS COURIR DERIER LES GROS SCORE	85.52387	43.83333	Control
8tvibhrk	bien comprendre les instructions et se concentrer sur la partie entrainement pour bien comprendre l experience	97.96217	63.66667	Control
j7o4zuva	priez dieu, car même en comprenant les règles, si vous n avez pas de chance, c est fini	97.08805	73.00000	Control
jdpddzwy	bien lire avant de ce lancer	89.31262	25.50000	Control
lxnap91b	prendre des pauses	87.72244	37.66667	Control
orbcofex	ne pas se précipiter	89.71518	57.00000	Control
4w0726be	Bien savoir les probas.	93.62210	32.50000	Control
bgkxjcv	Faites preuve de patience (et de chance...)	90.02617	37.66667	Control
ed5v7nok	Essayer a voir différentes stratégies en fonction du jeu	95.51103	36.66667	Control

<b>Participant</b>	<b>Advice</b>	<b>Perf</b>	<b>Explo_Quali</b>	<b>Group</b>
h56rxmif	Suivre son instinct	94.02494	59.33333	Control
mze8w1wb	Ne pas juste se baser uniquement sur le score A le plus élevé pour répondre aux questions et faire ses choix	95.06436	28.33333	Control
r78sqe2d	Lire avec attention	95.36259	77.50000	Control
d9i9ovbw	Confiance, ne pas trop réfléchir	94.70291	61.66667	Control
jl13do4x	Explorer les candidatures qui sont moins certaines d'être celles choisies	101.12679	66.50000	Control
x85nz122	Ne pas forcément regardé les scores de 10 car on les prends à chaque fois, préférer regarder les scores de 6	95.81363	56.83333	Control
1dorvkm3	Comprendre les consignes est l'élément clé.	92.75571	55.66667	Control
493q2yoi	Il faut choisir de révéler les scores B des candidats au score le plus proche de la médiane (5), car les deux expériences sont indépendantes, et identiques.	99.00263	58.66667	Control
6gqnc2xw	Bien penser au fait qu'on cherche à obtenir des informations pour ensuite statistiquement déduire qui pourrait être le meilleur candidat	95.60243	31.00000	Control
dk7pmxhr	de bien lire les consignes la première fois	99.18470	74.00000	Control
gkdettrf	Se faire confiance, fonder sa prise de décision majoritairement sur un calcul rationnel et en partie sur son intuition, sur sa capacité à faire un pari, par exemple si j'ai d'une part un score A de 5 avec un score B de 2 et d'autre part un score A de 4 avec un score B inconnu, je peux me permettre de prendre le risque de sélectionner le second candidat, en pariant sur le fait que sa compétence B sera d'un score supérieur à 2.	90.07260	34.66667	Control
o2ljnxyf	Etre attentif aux consignes et lire bien les instructions	93.53645	36.00000	Control

# General discussion

The general discussion serves as a synthesis and critical examination of the work presented throughout this thesis. It ties together the research objectives, experimental findings, and implications, addressing how the insights from this thesis advance our understanding of information search for set selection.

By taking a step back from the work carried out in this thesis, we leverage this discussion to critically assess the limitations of our results and their implications for future theoretical and practical research.

## Table of contents

Table of contents	188
I. Investigating information search before set selection	189
a. From recruiters to set selection	189
b. Research questions	189
II. Synthesis of the main results	189
a. Results by chapters	189
b. Specific contributions	190
III. Limitations and future research directions	191
a. Positioning within the “Rationality Wars”	191
b. Set Selection in traditional paradigms	194
c. Exploration and Selection Biases: relation to prior studies	196
d. Methodological Limitations	199
e. Other Limitations	200
IV. Perspectives on Real-Life Impacts	201
a. Personnel selection	201
b. Recommendation Algorithms	203
Bibliography	205

## I. Investigating information search before set selection

The general discussion begins with a recap of the thesis's initial context and the research questions addressed.

### a. From recruiters to set selection

The primary aim of this thesis was to understand how professionals, such as recruiters, make complex decisions in their daily activities. For instance, when recruiters engage in sourcing (searching for profiles on LinkedIn), they face an overwhelming number of possibilities for creating shortlists. This led us to question how individuals choose the information they wish to explore when tasked with selecting a set of options rather than a single option. As outlined in the introduction, numerous situations require set selection and demand specific, non-intuitive strategies, as opposed to choosing a single option.

### b. Research questions

In this thesis, we investigate how individuals conduct information searches when required to make set selections. Beginning with the premise that heuristics and biases often influence decision-making under bounded rationality, this thesis sought to uncover the cognitive mechanisms underlying exploration and selection strategies in set selection contexts. Key research questions included: What strategies do individuals use to make set selections? What information do they choose to explore to inform their selections? How do individuals manage to make a set selection in the face of uncertainty?

## II. Synthesis of the main results

In this section, we succinctly present the principal findings of the four chapters of the thesis. We describe the contributions of this work in relation to existing literature.

### a. Results by chapters

The findings of this thesis, derived from a series of experiments, underscore the challenges individuals encounter when selecting a set of options in decision-making tasks.

Chapter 1 demonstrated that most participants frequently exhibit two major biases: an exploration bias, where individuals disproportionately explore options perceived as favourites for selection, and a selection bias, where individuals are more likely to select options they have explored, even when these are not the most optimal choices.

Chapter 2 revealed that experts in recruitment exhibited the same exploration bias as non-experts. We also documented that these biases are only weakly correlated with general cognitive abilities and classical cognitive biases, which suggests that they cannot be reduced to e.g. confirmation bias.

Chapter 3 highlighted that participants tend to overestimate the effectiveness of their strategies, although their confidence aligns with the consistency of their chosen exploration strategies.

Finally, Chapter 4 provided evidence that participants can learn and adapt to optimal strategies when given a short training on information search in this task, though financial incentives failed to significantly influence behaviour or performance.

### b. Specific contributions

This thesis offers several distinct contributions to the literature on information search in decision-making, as outlined below.

First, we introduce a new theoretical question by examining set selection. This issue had not been addressed in the literature, despite its ecological relevance. Understanding how individuals resolve this type of problem is crucial, given its computational complexity and the counterintuitive strategies required compared to single-option selection.

Second, we developed a specific experimental paradigm to address our research question, rather than adapting an existing framework. The goal was to simulate a sourcing task in the laboratory—a common activity for recruiters. This paradigm was designed with an economist's approach, aiming at characterizing the information structure available to participants and at measuring the gap between observed and optimal behaviours. As demonstrated throughout the thesis, the paradigm is flexible (e.g., exploration set size, selection set size, framing, questionnaire versions). We hope it serves as a foundational framework for describing and studying information search in set-selection scenarios.

Third, this thesis adopts an interdisciplinary approach. The work primarily draws on literature from cognitive psychology and experimental economics. The task design follows typical experimental economics methodologies, such as incentivising choices (but also reported beliefs such as confidence in Chapter 3), using probabilistic gains and rewards based

on randomly chosen trials. We also relied on conceptual and experimental tools originating from cognitive psychology, such as the classification of cognitive biases (Ceschi et al., 2019), or the Cognitive Reflection Test (Brañas-Garza et al., 2019 for a review; Frederick, 2005) and the Berlin Numeracy Test (Cokely et al., 2012), now widely used in decision-making studies across various domains. At the crossroads of experimental economics and cognitive science is also the reliance on computational models and comparisons of participants' performances to that of optimal agents. Lastly, in Chapter 4, we simultaneously tested cognitive remediation methods (educational interventions) and economic interventions (financial incentives), allowing us to study their independent and combined effects.

### III. Limitations and future research directions

Having synthesised the key findings, we now turn to their limitations and implications for future research. In this section, we review the main limitations of this thesis, providing the framework within which our results can be interpreted and identifying areas requiring further investigation to draw broader conclusions. The discussion encompasses theoretical, methodological, and practical limitations in applying our findings to real-world situations. We also propose future research directions to address these limitations.

#### a. Positioning within the "Rationality Wars"

A significant debate in the fields of decision-making and reasoning, known as the "Rationality Wars" (Samuels et al., 2002; Sturm, 2012; Wallin, 2013), pits three main approaches against each other: logical rationality (Savage, 1954; Von Neumann & Morgenstern, 1944), the cognitive biases and heuristics program (Tversky & Kahneman, 1974), and ecological or bounded rationality (Gigerenzer & Goldstein, 1996).

While these approaches may seem opposed, they complement each other in explaining behaviour. This section attempts to position our results within "Rationality Wars".

#### **Logical Rationality**

Logical rationality defines the probabilistic rules that individuals should strictly follow to optimise decision-making and maximise their gains (Savage, 1954; Von Neumann & Morgenstern, 1944).

The optimal strategy described and used in this thesis is based on logical probability. However, we were unable to formalise it through a precise mathematical proof.

Consequently, calculating the expected gains of each alternative in each trial was impossible, and we relied on approximations and simulations.

Future research should aim to define a mathematical proof for the optimal strategy (if one exists). This would enable the precise calculation of the best possible expected gains and the exact deviation between optimal and observed performances. Such a model could serve as an educational tool for remediation by providing participants with a mathematical model of the optimal strategy. Additionally, it could facilitate the development of new experimental frameworks to which the same model could be applied.

### **Cognitive biases and heuristics program**

The biases and heuristics program demonstrates that individuals systematically deviate from logical reasoning (Kahneman & Tversky, 1973, 1979; Tversky & Kahneman, 1974, 1983). Exploration bias and selection bias are new examples that could be included in this framework.

In Chapter 2, we compared these biases with tasks measuring cognitive biases (e.g., confirmation bias, conjunction fallacy, and base-rate neglect) and cognitive abilities (e.g., the Cognitive Reflection Test and Berlin Numeracy Test). However, our results did not establish links between the biases described in this thesis and the standard cognitive bias tasks tested.

Further investigations are needed to determine whether cognitive factors common to these paradigms and the broader cognitive biases literature exist. Although we tested only a few relevant tasks, numerous others remain unexplored. As mentioned in Chapter 2, the work on taxonomies of cognitive biases didn't lead to a consensus yet (Berthet, Autissier, et al., 2022; Ceschi et al., 2019; Stanovich et al., 2008). New studies could determine whether common factors also apply to exploration and selection biases or whether these biases depend on distinct mechanisms, as suggested by the findings in Chapter 2.

### **Ecological rationality**

Ecological rationality posits that the heuristics individuals use are adaptive strategies suited to realistic and uncertain environments. These heuristics enable individuals to solve complex problems more efficiently and often with comparable accuracy than more complex methods that require participants to take into account more information (Gigerenzer & Goldstein, 1996; Gigerenzer & Todd, 1999).

Certain data from this thesis could be interpreted with the concept of ecological rationality. The strategy of exploring favourites is inexpensive to implement, as identifying the highest scores is straightforward. It is also the optimal approach for single selection. While it may be sub-optimal for ensemble selection, it is not entirely ineffective. The participants expressed satisfaction with it. In addition, participants demonstrated adaptability to task variations. In Chapter 1, their performance remained relatively stable, being lower when the task was objectively more challenging and slightly better when the task was easier. Their exploration shifted but not fully, suggesting that they adopted a trade-off between cognitive cost of full switch and benefits.

However, most data contradict the ecological rationality interpretation. First, the heuristic of exploring favourites may be an appropriate heuristic for a single selection but not the most effective for the set selection tasks in this thesis. Thus, its use does not reflect an ecologically rational choice. Second, in Chapter 4, some participants completed the task with a gain system that only rewarded perfect selections (e.g., selecting 5 out of the top 5 candidates earned €5, otherwise €0). Despite the suboptimality of their preferred heuristic under this system, no substantial changes in strategy were observed.

Future research assuming individuals maximise gains within cognitive and temporal constraints could better elucidate why participants persistently use the heuristic of exploring favourites. One hypothesis is that cognitive blinders prevent them from recognising alternative strategies. When these blinders are removed, such as through educational interventions, participants can adopt alternative strategies and perform better on task variations.

If individuals possess ecological rationality, investigating their exploration strategies' motivations may require understanding how they anticipate the selection phase. Their exploration strategy may be driven by anticipating a selection strategy different from the mathematically optimal one. To test this hypothesis, participants could be given a predefined exploration and their selection analysed. If participants perform better with the favourite exploration strategy than with the mathematically optimal exploration strategy, this suggests their exploration heuristic is better suited to their anticipated selection strategy.

## b. Set Selection in traditional paradigms

To better understand the effects of transitioning from single to set selection, further research using alternative paradigms will be necessary. With our paradigm, we observe that individuals tend to favour strategies that would be optimal for single selection. It raises the question of whether this phenomenon also occurs in other paradigms. If individuals consistently apply strategies typically used for single selection, it could indicate a general difficulty in adapting to set selection, rather than a challenge specific to our paradigm. To study set selection and compare participants' strategies between single and set selection, here we propose modified version of two paradigms introduced earlier, namely the exploration-exploitation and the sequential evidence accumulation paradigm (see section IV.b and IV.c in our General Introduction). These paradigms are particularly relevant for studying set selection as they replicate realistic decision-making scenarios.

### **Exploration–Exploitation Paradigm**

To recap, the exploration–exploitation paradigm describes the dilemma between exploring new options to gather more information and exploiting known options to maximise immediate gains. Multi-armed bandit tasks simulate this paradigm by asking participants to choose among several "slot machines" (bandits), each with unknown reward probabilities, requiring them to balance discovering the best options (exploration) with maximising rewards (exploitation).

This paradigm is particularly well-suited for adaptation from single to set selection, as individuals can explore/exploit multiple options simultaneously rather than one at a time. The design doesn't need many adjustments to make it understandable. The participant faces 10 one-armed bandits. At each trial, he must choose 3 different arms to activate. The value of each bandit is dynamic: the value changes with each trial, whether or not the bandit has been activated. On each trial, activating the best bandit is considered as exploitation. The other arms are considered explorations.

The set selection version may be representative of real-life situations. For instance, a researcher might simultaneously pursue multiple projects: a primary focus as a "safe bet" (exploitation) alongside several exploratory collaborations (exploration). Similarly, businesses often balance core activities with exploratory ventures; for example, a consulting firm (exploitation) may experiment with coaching programmes (exploration).

In this context, one could investigate how individuals allocate exploration to maintain their primary source of exploitation while adapting to substantial variations in the value of alternatives. This paradigm lacks a calculable optimal solution when there is a lot of bandits, and the values of different bandits fluctuate over time, regardless of whether they have been explored (Cohen et al., 2007; Daw et al., 2006). Consequently, there would also be no theoretical optimal performance to calculate for a scenario involving multiple actions per trial.

What can we expect then in this modified paradigm? In the classical version, the behaviour of individuals can be described by a softmax rule, i.e. they maximise their gains by exploiting the most profitable bandit, but not perfectly (Cohen et al., 2007). One hypothesis is that the increase in information to integrate could amplify noise in information processing, making it harder to maximise gains. Another hypothesis could be that participants might learn more efficiently which arm-bandit are the bests. In the original version, with only one action per trial, the participant is forced to abandon exploiting the best slot machine. In the modified version involving ensemble selection, they could simultaneously exploit the most rewarding bandit while exploring by giving up, for instance, the 3<sup>rd</sup> most rewarding bandit. In this sense, this adapted version of the exploration-exploitation dilemma differs from our paradigm of set selection - set exploration. For exploration-exploitation, the optimal strategy is to retain the most rewarding bandit in both single and set selection scenarios.

### **Sequential Evidence Accumulation Paradigm**

The sequential evidence accumulation paradigm involves paying to receive pieces of information one at a time until the individual decides to stop gathering information and choose from the available options. As mentioned earlier, a horse race betting scenario provides an illustrative example. Each piece of information provided to participants corresponds to a betting tip with an associated reliability score, such as betting on horse Pegasus with 75% reliability. The participant doesn't choose about what horse one receive information (Hausmann & Läge, 2008).

In horse race betting, many formats involve set selection. For instance, the quinté+ requires predicting the top five horses in the correct order, and the tiercé involves identifying the top three in order. It would thus make sense to adapt the evidence accumulation task into one suitable for set selection.

If participants were required to predict the top five finishers (i.e., the quinté) instead of the winning horse, only the selection stage needs adaptation. During the accumulation phase,

each piece of information concerns one horse, randomly chosen. Participants would then stop accumulating information when they feel ready to provide a quinté prediction and not only the winner.

This adaptation would allow for an analysis of whether confidence plays the same role. Hausmann & Läge's study (2008) demonstrates each individual try to reach a individual desired level of confidence (DLC) before committing to a decision. Moreover, the concept of confidence itself could be further explored by examining how confidence in each individual option included in the quinté (local confidence) translates into confidence in the entire quinté prediction (global confidence). Does each option need to meet the DLC threshold to be included in the quinté, or is the required threshold a combination of confidence scores across all options? Answering this question would provide deeper insights into the connection between single-option and set-based decision-making within a sequential evidence accumulation process.

### c. Exploration and Selection Biases: relation to prior studies

We have described two major biases present in the exploration and selection phases of our paradigm. Further studies are necessary to understand whether these biases differ from those introduced earlier in the thesis. We wonder if exploration and selection bias can be derived from the same mechanisms as the concept we describe below.

#### **Confirmation bias: a distinct mechanism**

In Chapter 2, we tested whether confirmation bias could explain exploration and selection biases, but the results were inconclusive. The logical  $P \rightarrow Q$  hypothesis framework commonly associated with confirmation bias tasks (Berthet, Teovanovic, et al., 2022; Wason, 1960, 1968) does not align with our paradigm, leading us to dismiss this hypothesis.

#### **Pavlovian Biases**

Pavlovian biases describe tendencies to approach options associated with rewards. These biases manifest as "Sampling the Favourite" (exploring the favourite option), "Positive Evidence Approach" (selecting the option when new information confirms my preference or this option) and "Rejecting the Unsampled Options" (not selecting an option that you have refused to explore) (Hunt et al., 2016). The "Positive Evidence Approach" bias cannot be identified in our paradigm because it requires sequential accumulation of at least two pieces of information, whereas our paradigm includes only one step of information exploration. However, the other two biases closely resemble exploration bias towards favourites and

selection bias towards explored options reported in our data. Selection bias can be framed as a bias against unexplored options.

Specific data reinforce similarities between the biases described by Hunt et al. (2016) and those observed in our thesis. In Chapter 1, participants tended to reject unexplored options more frequently when tasked with identifying the top five candidates compared to eliminating the bottom five. This aligns with the "Rejecting the Unsampled Options" bias.

In Chapter 4, despite the educational intervention improving exploration strategies, participants still exhibited an exploration bias towards favourites. This bias was present across all experiments in the thesis, except in Experiment 3 (Chapter 1), where the average bias was neutral. The exception was experiment 3 in Chap. 1, in which the task was to eliminate the 5 worst candidates. The bias was zero on average. This neutrality stemmed from some participants favouring action-related priorities (e.g., rejecting the bottom candidates), while others prioritised favourites based on rankings (e.g., selecting candidates with the best A scores).

The hypothesis of a Pavlovian source for these biases warrants further investigation. Modifications to our paradigm, such as allowing participants to accumulate more information about candidates, could enable testing for the "Positive Evidence Approach" bias.

### **Choice bracketing**

Choice bracketing refers to the level of perspective with which a decision is made. Decisions are improved when taken with a global objective - broad bracketing - rather than a local one - narrow bracketing. For instance, Read et al. (2000) examined financial decision-making by asking participants to allocate a fictional budget over a given period. Narrow bracketing involved segmenting decisions with daily allocations, while broad bracketing entailed optimising the allocation across the entire period.

Our task seems to use broad bracketing, as individuals can consider their explorations simultaneously, adopting a global perspective on the entire exploration phase and then on the entire selection phase. They therefore can employ a broad perspective at each stage. Narrow bracketing, on the other hand, would involve asking participants to conduct a single exploration and then decide whether to select or reject that option. The participant would thus repeat the cycle of exploration and subsequent selection or rejection five times in succession.

If participants do not perform better or worse in the narrow bracketing version of our paradigm, it is likely that choice bracketing has no effect on participants in our task. However, if their performance is worse compared to the original version, this would suggest that broad bracketing serves to minimise performance loss relative to the optimal strategy.

Conversely, if participants perform better in the narrow bracketing version than in the original task, it could indicate that our task aligns with the rare cases where narrow bracketing is advantageous. Read et al. (2000) cites the example of taxi drivers who set daily mileage goals rather than weekly ones to focus on an achievable target. An intuitive hypothesis is that participants may better understand which information is useful to explore in our task if they are required to explore and decide on options one by one. They might then realise that they always retain their favourite option after exploring it, regardless of its score B.

### **Failure of Contingent Thinking**

In Chapter 4, we interpreted the success of the educational intervention as resolving individuals' failure of contingent thinking (FCT). FCT posits that individuals make suboptimal choices when unable to consider all aspects of a task and anticipate possible scenarios (Martínez-Marquina et al., 2019). For example, in our task, the implicit objective was to explore options that would be most informative for identifying candidates within vs. outside the top five. The educational intervention explicitly clarified which information was useful and why, helping participants anticipate the selection phase when they were deciding which options to explore.

For FCT, Esponda & Vespa (2014) used a committee voting task where participants vote alongside two computers: one always votes blue, and the other votes based on the actual outcome. The challenge is to vote optimally considering the pivotal role of the human voter. Suboptimal behaviour arises when participants neglect the importance of being decisive. Replicating such tasks could provide insights into the cognitive mechanisms behind contingent reasoning in set selection.

### **Uncertainty aversion and Risk aversion**

Selection bias towards explored options was not a central focus of this thesis. One could speculate that potential explanations for this bias include risk aversion and uncertainty aversion.

However, some of our empirical data suggests risk aversion is not at stake here. In Chapter 2, we tested risk aversion using a task requiring participants to make 10 binary choices between a "safe" lottery with low variance in rewards and a "risky" lottery with higher variance (Holt & Laury, 2002). Results showed no link between the number of risky choices and selection bias.

An alternative explanation could be uncertainty aversion, which reflects a tendency to avoid situations where outcome probabilities are unknown or ambiguous, favouring options with well-defined probabilities instead (Ellsberg, 1961). The distinction between risk aversion and uncertainty aversion lies in whether probabilities are known or unknown. For instance, in a task involving an urn with 50 blue and 50 yellow balls, participants face risk when choosing between a €50 guaranteed gain or a lottery offering €100 if a blue ball is drawn. By contrast, uncertainty arises when choosing between similar options without knowing the ball distribution.

To draw conclusions about the relationship between selection bias and risk or uncertainty aversion, additional tasks measuring these constructs would need to be tested. It is unclear whether selection in our paradigm relates more to risk or uncertainty. It is unclear if our paradigm relates to risk or to uncertainty given the complexity of probabilities computing. Comparing inter-individual variations in selection bias with tasks assessing uncertainty aversion could provide deeper insights into the drivers of set selection.

### d. Methodological Limitations

This thesis predominantly relies on laboratory experiments, which, while invaluable for controlled investigation, may limit the ecological validity of its findings. The tasks designed, such as the two-stage exploration and selection paradigm, simplify real-world complexities to focus on specific cognitive mechanisms. These simplifications included for instance the choice of no correlation between scores A and scores B, the fact that participants received the information on B scores only after choosing all the explored options (alternatively, one could explore one option at a time) and the use of a fixed number of options to explore. Such simplifications may not fully capture the dynamics of decision-making in practical settings, such as recruitment or academic selection processes.

Moreover, the participant samples were mainly composed of French adults from a single experimental pool in Paris (with the exception of the expert sample). All experiments were

conducted online. This limits the generalisability of the findings across diverse populations and contexts. Although we note that results showed no significant effects of expertise, age, education level, or gender, we acknowledge that future research should aim to replicate these findings with more diverse populations and in field settings to evaluate the robustness of the identified biases and heuristics.

### e. Other Limitations

Factors such as personality traits and non-instrumental value of information were not explicitly considered in our experimental work, even though these often play crucial roles in real-world decision-making and could explain the inter-individual variations observed in biases.

#### **Inter-Individual Differences**

Our results highlight individual differences that could be explored in future research. Through the experiments, we observe that some participants exhibit neither exploration nor selection bias. Interestingly, some participants even display opposite biases: an exploration bias towards outsiders and a selection bias towards unexplored options. Inter individual differences are also evident in participants' assessments of the quality of their exploration: some overestimate their performance, while others underestimate it (Chapter 3). Moreover, a portion of participants failed to improve their exploration strategy despite the educational intervention. Conversely, other participants required no intervention and demonstrated an effective strategy for both the original task and its variations (Chapter 4).

Incorporating psychological and cognitive measures, such as personality assessments or neuropsychological tests, could help examine how personality traits influence exploration and selection behaviours. Recent studies show inter-individual differences in cognitive bias expression could be related to personality traits (Ahmad, 2020; Kumar et al., 2023; Singh et al., 2023). For example, extraversion is strongly linked to overconfidence bias and confirmation bias (Ahmad, 2020; Kumar et al., 2023). In contrast, neuroticism leads to risk-averse strategies (Singh et al., 2023). Exploring links with personality traits could help us to describe behaviours more fully in our paradigm.

#### **Non-instrumental value of information**

The optimal strategy assumes a specific utility framework, which may not align with the individual goals of participants. For instance, some participants might prioritise the non-instrumental value of information—the intrinsic satisfaction or confidence gained from

acquiring information—over maximising expected gains (Eliaz & Schotter, 2010; Kobayashi et al., 2019). In such cases, participants may focus on enhancing their confidence in pre-decisions (i.e., the choices they anticipate making prior to exploration) rather than pursuing purely instrumental objectives.

To investigate if the exploration bias could be applied to a task in which there is no instrumental value of information. We could make a task in which exploration stage take place after the selection stage. For example, participants could make their selection based solely on scores A, followed by an exploration phase to gather information about scores B for five options. At the end of the trial, we would provide them with feedback on the quality of their selection, ensuring that the exploration phase is not essential for determining the quality of their decision.

## IV. Perspectives on Real-Life Impacts

Based on the findings of this thesis, we can envision the potential impact of its insights on real-world situations, such as recruitment or recommendation algorithms. We revisit the recruitment domain, which served as a contextual backdrop for this research, and discuss the potential implications in the field of recommendation algorithms. We believe that biases in set selection have yet to be recognised as significant issues, despite their potentially large effects.

We describe the impact these biases may have when combined with other types of biases (stereotypes, algorithms biases), along with potential strategies to address them. These implications highlight the practical relevance of this thesis to real-world challenges, offering actionable strategies to enhance decision quality in complex environments. By tackling the biases inherent in set selection, organisations and professionals can promote fairer and more effective decision-making systems.

### a. Personnel selection

#### **Resume screening**

Numerous studies focus on discrimination in resume screening, highlighting its role in reducing candidates' chances of securing a job interview and creating inequalities in access to employment (Bertrand & Duflo, 2017; Eying, 2022; Kroll et al., 2021; Neumark, 2018; Pager et al., 2009). This first stage in the recruitment process generally involves selecting sets of

potential candidates from a very large pool. In this respect, it corresponds to the paradigm presented in this thesis.

Encouraging solutions have been proposed to reduce discrimination at this stage. Some studies have shown the positive effects of interventions to help recruiters reduce their stereotypes, although the long-term effects have been mixed (Devine et al., 2012; Forscher et al., 2017). One solution that has been tested is the introduction of anonymous CVs. Once again, the effects are mixed, with positive effects particularly for women (Åslund & Skans, 2012), but unexpected negative effects for people with less conventional backgrounds, partly because of previous discrimination (Behaghel et al., 2015). Overall, this first sorting stage, which is supposed to lead to a selection of all candidates, has not yet been debiased and is generating major inequalities in access to employment (Petit et al., 2013).

### **Interactions with exploration bias and consequences**

Biases such as stereotypes in recruitment process (Bosak & Sczesny, 2011; Koch et al., 2015; Sczesny et al., 2004) can interact with exploration biases. Stereotypes influence the recruitment process even before candidates are evaluated, acting as an initial filter for recruiters. In our paradigm, stereotypes function similarly to the “score A”. Instead of receiving an objective initial score, recruiters rely on a biased subjective assessment. As a result, stereotypes determine which candidates are seen as favourites for selection.

When stereotypes drive the identification of favourites, individuals subject to these biases are less likely to have their applications explored (Bartoš et al., 2016). This compounds the issue, as their resumes are also evaluated less favourably when reviewed. Therefore, the harm can be twofold: fewer opportunities to be considered for the selection set and lower chances of being evaluated positively if considered. Overall, this describes the many reasons why people who are victims of stereotyping are rejected, even at stages where a whole range of candidates are still in the run.

Strategies can be devised to mitigate these effects. For instance, the optimal exploration strategy outlined in this thesis may not be feasible in certain contexts. As discussed in Chapter 1, student selection procedures for academic programs often preclude selecting a candidate with an average academic record but an excellent interview performance at the expense of another candidate with a slightly better academic record who was not shortlisted for an interview.

Other solutions could address the combined effects of these practices. In France, quotas cannot legally be applied to criteria such as ethnicity or age, unlike gender or disability. A more general solution might involve a right to a second chance. At each stage of the recruitment process, randomly selected excluded candidates could be reintroduced into subsequent stages. This approach would allow for recruiter errors to be rectified and offer candidates another opportunity to be considered. Such a solution would complement other efforts to improve recruitment processes, including recruiter training, the use of structured interviews, or predictive validity tests.

### b. Recommendation Algorithms

Recommendation algorithms are ubiquitous in our lives, integrated into nearly every digital tool. Their main purpose is to provide us with relevant content to help us make the right decisions. On social media platforms such as Facebook, Instagram, X, and TikTok, or streaming services like YouTube, Netflix, and Spotify, they shape the content we see, tailoring it to our preferences. On e-commerce sites like Amazon or Shein, they suggest products we are most likely to purchase. News websites use them to recommend articles related to the one just read. Even in the scientific domain, researchers rely heavily on the recommendations of Google Scholar's algorithm following keyword searches.

#### **Algorithmic Biases**

Algorithmic biases, defined as systematic distortions in the outcomes produced by automated systems, influence how users perceive and interact with content. For instance, popularity bias in recommendation systems amplifies the visibility of already popular items at the expense of less visible ones, creating a "rich get richer" and "poor get poorer" effect (Abdollahpouri, 2019). Another example is provided by (Cho et al., 2020), who demonstrate how YouTube's algorithms exacerbate political polarisation by recommending videos aligned with the user's ideological preferences. These biases are not confined to data but are also embedded within models (Kordzadeh & Ghasemaghaei, 2022).

While several solutions have been proposed to mitigate algorithmic biases, their effectiveness remains mixed depending on the context and implementation. For instance, techniques such as reweighting data or integrating fairness metrics into algorithms, as suggested by Chen et al. (2023), have shown potential in reducing popularity bias and improving exposure diversity. However, these approaches often face trade-offs, such as reduced overall accuracy or increased computational complexity (Abdollahpouri & Mansoury, 2020). User-driven solutions, like the interactive tools tested by Millecamp et al.

(2018) on Spotify, empower individuals to control recommendations, yet their success depends heavily on user engagement and understanding. Transparency and audit practices have been effective in identifying systemic issues (Kordzadeh & Ghasemaghaei, 2022), but implementing these measures at scale remains a challenge, highlighting the need for more robust, context-specific interventions.

### **Interactions with exploration bias and consequences**

While human-machine interactions are fundamental to the use of recommendation algorithms, it is striking how their design often disregards the complexity of human behaviour. Popularity bias serves as a clear example of how human preferences for favourites are mirrored in recommendation systems. By embedding a human bias into a recommendation algorithm, the bias is effectively amplified. Consequently, individuals are exposed solely to popular or favoured content, narrowing their exploration to these limited options before ultimately selecting content they are most familiar with. This polarisation towards popular content arises not just from the algorithm itself but from the interaction between a biased individual and an equally biased algorithm.

Recommendation algorithms can be designed to actively reduce our biases. Current proposals to address algorithmic biases often aim to eliminate bias entirely, striving for neutrality to enhance the effectiveness of recommendation systems. However, more can be done. By adopting an interdisciplinary approach that combines information sciences and behavioural sciences, algorithms could be intentionally designed to counteract our biases. For example, algorithms could purposefully recommend “outsiders” or unfamiliar content. Such a counterintuitive approach would require these recommendations to be presented transparently (to avoid misleading users) but also attractively. In practice, algorithms could frame content as opportunities to be a film buff when searching for a new series on Netflix, suggest articles to think “out of the box” when browsing press articles, or promote becoming an “early fan” of an emerging artist through unexpected recommendations on Spotify.

## Bibliography

1. Abdollahpouri, H. (2019). Popularity Bias in Ranking and Recommendation. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 529–530. <https://doi.org/10.1145/3306618.3314309>
2. Abdollahpouri, H., & Mansoury, M. (2020). *Multi-sided Exposure Bias in Recommendation* (arXiv:2006.15772). arXiv. <https://doi.org/10.48550/arXiv.2006.15772>
3. Ahmad, F. (2020). Personality traits as predictor of cognitive biases: Moderating role of risk-attitude. *Qualitative Research in Financial Markets*, 12(4), 465–484. <https://doi.org/10.1108/QRFM-10-2019-0123>
4. Åslund, O., & Skans, O. N. (2012). Do Anonymous Job Application Procedures Level the Playing Field? *ILR Review*, 65(1), 82–107. <https://doi.org/10.1177/001979391206500105>
5. Bartoš, V., Bauer, M., Chytilová, J., & Matějka, F. (2016). Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition. *American Economic Review*, 106(6), 1437–1475. <https://doi.org/10.1257/aer.20140571>
6. Behaghel, L., Crépon, B., & Le Barbanchon, T. (2015). Unintended Effects of Anonymous Résumés. *American Economic Journal: Applied Economics*, 7(3), 1–27. <https://doi.org/10.1257/app.20140185>
7. Berthet, V., Autissier, D., & de Gardelle, V. (2022). Individual differences in decision-making: A test of a one-factor model of rationality. *Personality and Individual Differences*, 189, 111485. <https://doi.org/10.1016/j.paid.2021.111485>
8. Berthet, V., Teovanovic, P., & de Gardelle, V. (2022). *Confirmation bias in hypothesis testing: A unitary phenomenon?* [Preprint]. Open Science Framework. <https://doi.org/10.31219/osf.io/wjkr5>
9. Bertrand, M., & Duflo, E. (2017). Field Experiments on Discrimination. In A. V. Banerjee & E. Duflo (Eds.), *Handbook of Economic Field Experiments* (Vol. 1, pp. 309–393). North-Holland. <https://doi.org/10.1016/bs.hefe.2016.08.004>
10. Bosak, J., & Sczesny, S. (2011). Gender Bias in Leader Selection? Evidence from a Hiring Simulation Study. *Sex Roles*, 9.
11. Brañas-Garza, P., Kujal, P., & Lenkei, B. (2019). Cognitive reflection test: Whom, how, when. *Journal of Behavioral and Experimental Economics*, 82, 101455. <https://doi.org/10.1016/j.socec.2019.101455>
12. Ceschi, A., Costantini, A., Sartori, R., Weller, J., & Di Fabio, A. (2019). Dimensions of decision-making: An evidence-based classification of heuristics and biases. *Personality and Individual Differences*, 146, 188–200. <https://doi.org/10.1016/j.paid.2018.07.033>
13. Chen, J., Dong, H., Wang, X., Feng, F., Wang, M., & He, X. (2023). Bias and Debias in Recommender System: A Survey and Future Directions. *ACM Trans. Inf. Syst.*, 41(3), 67:1–67:39. <https://doi.org/10.1145/3564284>

14. Cho, J., Ahmed, S., Hilbert, M., Liu, B., & Luu, J. (2020). Do Search Algorithms Endanger Democracy? An Experimental Investigation of Algorithm Effects on Political Polarization. *Journal of Broadcasting & Electronic Media*.  
<https://www.tandfonline.com/doi/abs/10.1080/08838151.2020.1757365>
15. Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933–942. <https://doi.org/10.1098/rstb.2007.2098>
16. Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring Risk Literacy: The Berlin Numeracy Test. *Judgment and Decision Making*, 7(1), 25–47.  
<https://doi.org/10.1017/S1930297500001819>
17. Daw, N. D., O’Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876–879.  
<https://doi.org/10.1038/nature04766>
18. Devine, P. G., Forscher, P. S., Austin, A. J., & Cox, W. T. L. (2012). Long-term reduction in implicit race bias: A prejudice habit-breaking intervention. *Journal of Experimental Social Psychology*, 48(6), 1267–1278. <https://doi.org/10.1016/j.jesp.2012.06.003>
19. Eliaz, K., & Schotter, A. (2010). Paying for confidence: An experimental study of the demand for non-instrumental information. *Games and Economic Behavior*, 70(2), 304–324.  
<https://doi.org/10.1016/j.geb.2010.01.006>
20. Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms\*. *The Quarterly Journal of Economics*, 75(4), 643–669. <https://doi.org/10.2307/1884324>
21. Esponda, I., & Vespa, E. (2014). Hypothetical Thinking and Information Extraction in the Laboratory. *American Economic Journal: Microeconomics*, 6(4), 180–202.  
<https://doi.org/10.1257/mic.6.4.180>
22. Eyting, M. (2022). Why do we Discriminate? The Role of Motivated Reasoning. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4210315>
23. Forscher, P. S., Mitamura, C., Dix, E. L., Cox, W. T. L., & Devine, P. G. (2017). Breaking the prejudice habit: Mechanisms, timecourse, and longevity. *Journal of Experimental Social Psychology*, 72, 133–146. <https://doi.org/10.1016/j.jesp.2017.04.009>
24. Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, 19(4), 25–42. <https://doi.org/10.1257/089533005775196732>
25. Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650–669. <https://doi.org/10.1037/0033-295X.103.4.650>
26. Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart* (pp. 3–34). Oxford University Press.

27. Hausmann, D., & Läge, D. (2008). Sequential evidence accumulation in decision making: The individual desired level of confidence can explain the extent of information acquisition. *Judgment and Decision Making*, 3(3), 229–243. <https://doi.org/10.1017/S1930297500002436>
28. Holt, C. A., & Laury, S. K. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 92(5), 1644–1655. <https://doi.org/10.1257/000282802762024700>
29. Hunt, L. T., Rutledge, R. B., Malalasekera, W. M. N., Kennerley, S. W., & Dolan, R. J. (2016). Approach-Induced Biases in Human Information Sampling. *PLOS Biology*, 14(11), e2000638. <https://doi.org/10.1371/journal.pbio.2000638>
30. Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237.
31. Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263–291. <https://doi.org/10.2307/1914185>
32. Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature Human Behaviour*, 3(6), 587–595. <https://doi.org/10.1038/s41562-019-0589-3>
33. Koch, A. J., D’Mello, S. D., & Sackett, P. R. (2015). A meta-analysis of gender stereotypes and bias in experimental simulations of employment decision making. *Journal of Applied Psychology*, 100(1), 128–161. <https://doi.org/10.1037/a0036734>
34. Kordzadeh, N., & Ghasemaghahi, M. (2022). Algorithmic bias: Review, synthesis, and future research directions. *European Journal of Information Systems*, 31(3), 388–409. <https://doi.org/10.1080/0960085X.2021.1927212>
35. Kroll, E., Veit, S., & Ziegler, M. (2021). The Discriminatory Potential of Modern Recruitment Trends—A Mixed-Method Study From Germany. *Frontiers in Psychology*, 12. <https://www.frontiersin.org/articles/10.3389/fpsyg.2021.634376>
36. Kumar, V., Dudani, R., & K., L. (2023). The big five personality traits and psychological biases: An exploratory study. *Current Psychology*, 42(8), 6587–6597. <https://doi.org/10.1007/s12144-021-01999-8>
37. Martínez-Marquina, A., Niederle, M., & Vespa, E. (2019). Failures in Contingent Reasoning: The Role of Uncertainty. *American Economic Review*, 109(10), 3437–3474. <https://doi.org/10.1257/aer.20171764>
38. Millicamp, M., Htun, N. N., Jin, Y., & Verbert, K. (2018). Controlling Spotify Recommendations: Effects of Personal Characteristics on Music Recommender User Interfaces. *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*, 101–109. <https://doi.org/10.1145/3209219.3209223>
39. Neumark, D. (2018). Experimental Research on Labor Market Discrimination. *Journal of Economic Literature*, 56(3), 799–866. <https://doi.org/10.1257/jel.20161309>

40. Pager, D., Bonikowski, B., & Western, B. (2009). Discrimination in a Low-Wage Labor Market: A Field Experiment. *American Sociological Review*, 74(5), 777–799.  
<https://doi.org/10.1177/000312240907400505>
41. Petit, P., Duguet, E., L'Horty, Y., du Parquet, L., & Sari, F. (2013). Discrimination à l'embauche: Les effets du genre et de l'origine se cumulent-ils systématiquement ? *Economie et statistique*, 464(1), 141–153. <https://doi.org/10.3406/estat.2013.10234>
42. Read, D., Loewenstein, G., Rabin, M., Keren, G., & Laibson, D. (2000). Choice Bracketing. In B. Fischhoff & C. F. Manski (Eds.), *Elicitation of Preferences* (pp. 171–202). Springer Netherlands.  
[https://doi.org/10.1007/978-94-017-1406-8\\_7](https://doi.org/10.1007/978-94-017-1406-8_7)
43. Samuels, R., Stich, S., & Bishop, M. (2002). Ending the Rationality Wars How to Make Disputes about Human Rationality Disappear. In R. Elio (Ed.), *Common Sense, Reasoning, and Rationality* (1st ed., pp. 236–268). Oxford University Press New York.  
<https://doi.org/10.1093/0195147669.003.0011>
44. Savage, L. J. (1954). *The foundations of statistics* (pp. xv, 294). John Wiley & Sons.
45. Sczesny, S., Bosak, J., Neff, D., & Schyns, B. (2004). Gender Stereotypes and the Attribution of Leadership Traits: A Cross-Cultural Comparison. *Sex Roles*, 51(11–12), 631–645.  
<https://doi.org/10.1007/s11199-004-0715-0>
46. Singh, Y., Adil, Mohd., & Haque, S. M. I. (2023). Personality traits and behaviour biases: The moderating role of risk-tolerance. *Quality & Quantity*, 57(4), 3549–3573.  
<https://doi.org/10.1007/s11135-022-01516-4>
47. Stanovich, K. E., Toplak, M. E., & West, R. F. (2008). The Development of Rational Thought: A Taxonomy of Heuristics and Biases. In R. V. Kail (Ed.), *Advances in Child Development and Behavior* (Vol. 36, pp. 251–285). JAI. [https://doi.org/10.1016/S0065-2407\(08\)00006-2](https://doi.org/10.1016/S0065-2407(08)00006-2)
48. Sturm, T. (2012). The “Rationality Wars” in Psychology: Where They Are and Where They Could Go. *Inquiry*, 55(1), 66–81. <https://doi.org/10.1080/0020174X.2012.643628>
49. Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
50. Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315.  
<https://doi.org/10.1037/0033-295X.90.4.293>
51. Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior* (pp. xviii, 625). Princeton University Press.
52. Wallin, A. (2013). A peace treaty for the rationality wars? External validity and its relation to normative and descriptive theories of rationality. *Theory & Psychology*, 23(4), 458–478.  
<https://doi.org/10.1177/0959354313489369>

## General discussion

53. Wason, P. C. (1960). On the Failure to Eliminate Hypotheses in a Conceptual Task. *Quarterly Journal of Experimental Psychology*, 12(3), 129–140. <https://doi.org/10.1080/17470216008416717>
54. Wason, P. C. (1968). Reasoning about a Rule. *Quarterly Journal of Experimental Psychology*, 20(3), 273–281. <https://doi.org/10.1080/14640746808400161>

# Résumé substantiel en français

## Introduction

« Dans un monde riche en informations, la profusion d'informations signifie une pénurie d'autre chose : une rareté de ce que l'information consomme. Ce que l'information consomme est assez évident : elle consomme l'attention de ses récepteurs. Ainsi, une richesse d'informations crée une pauvreté d'attention et la nécessité d'allouer cette attention de manière efficace parmi l'abondance des sources d'information susceptibles de la consommer. »

— Herbert A. Simon

*'Designing Organizations for an Information-Rich World'* in Martin Greenberger (ed.) *Computers, Communications, and the Public Interest* (1971)

Environ 252 000 nouveaux sites web sont créés chaque jour (Drezner & Edigbe, 2024). According to Statista, 347 billion emails are sent globally each day. Selon Statista, 347 milliards d'e-mails sont envoyés quotidiennement à travers le monde. Une personne reçoit en moyenne 43 e-mails par jour (Statista, 2024). Comme l'a souligné Herbert Simon, la capacité à allouer efficacement son attention parmi une pléthore d'informations est cruciale. Dans cette thèse, nous nous concentrons sur l'optimisation de la recherche d'informations pour prendre une décision complexe impliquant la sélection simultanée de plusieurs options, et non d'une seule.

# Table des matières

Introduction.....	210
Table des matières.....	211
Origine du projet de thèse : prise de décision chez les recruteurs-----	212
Sélection unique vs sélection d'un ensemble : quelle différence cela fait-il ? -----	214
Méthode Générale : Développement d'un paradigme expérimental pour l'étude de la recherche d'informations dans la sélection d'ensemble .....	217
Paradigme d'exploration et de sélection d'ensemble-----	218
Caractéristiques clés du paradigme -----	218
Méthodologie expérimentale (du Chapitre 1) .....	220
Stimuli et procédure -----	220
Modèle et mesures -----	221
Chapitre 1 : Biais d'exploration en faveur des options favorites dans la sélection d'ensemble .....	223
Chapitre 2 : Facteurs cognitifs influençant les biais dans la recherche d'informations .....	224
Chapitre 3 : Jugements métacognitifs sur les stratégies d'exploration pour la sélection d'ensemble.....	226
Chapitre 4 : Apprentissage et adaptation des stratégies optimales.....	227
Discussion .....	228
Contributions spécifiques -----	228
Limites méthodologiques -----	229
Autres Limitations -----	230
Perspectives sur les impacts dans la vie réelle.....	231
Sélection du personnel-----	232
Algorithmes de recommandation -----	233
Bibliographie .....	236

## Origine du projet de thèse : prise de décision chez les recruteurs

Pour commencer, nous décrivons le cadre initial de cette thèse, qui se concentre sur les processus de prise de décision des recruteurs. Nous expliquons comment nous sommes partis d'une problématique liée au recrutement, connue sous le nom de "sourcing", pour étudier la sélection d'ensembles dans un environnement contrôlé de laboratoire à travers cette thèse.

### **Le sourcing : la recherche de talents sur LinkedIn**

Le sourcing est une pratique de recrutement qui consiste à contacter directement des candidats plutôt que d'attendre des réponses à une offre d'emploi publiée. Cette pratique fait partie des techniques de recrutement en pleine expansion, notamment grâce aux réseaux sociaux comme LinkedIn, qui permettent aux recruteurs d'approcher des candidats passifs, c'est-à-dire des individus déjà employés mais ouverts à de nouvelles opportunités (Hosain & Liu, 2020). Selon une étude menée par *LinkedIn* en 2015, 75 % des travailleurs se considéraient comme des candidats passifs.

L'un des premiers défis pour les recruteurs est de pouvoir filtrer les profils pertinents avant d'établir un contact. LinkedIn affirme compter plus d'un milliard d'utilisateurs dans le monde en 2024, dont 30 millions en France. En seulement deux ans, le nombre d'utilisateurs français aurait augmenté de 20 %. Une recherche simple pour "Data Engineer" basé en France donne 46 000 profils. Examiner tous ces profils est clairement impossible, et même passer en revue les détails des 10 premiers affichés dans les résultats de recherche prendrait trop de temps.

Un autre défi est d'identifier les informations pertinentes à explorer. Chaque profil comporte plusieurs sections (expérience, formation, activités sur le site, etc.), chacune pouvant être très détaillée. Chaque utilisateur complète son profil à des niveaux de détail variables. Un recruteur doit trier toutes les informations disponibles pour déterminer si un profil mérite d'être pris en compte dans le processus de recrutement. En général, il est impossible pour un recruteur d'explorer toutes les informations pertinentes pour chaque profil susceptible d'être intéressant. L'objectif est donc d'identifier les défis auxquels les recruteurs peuvent être confrontés dans cette pratique, qui devient une partie de plus en plus importante de leur rôle. Selon l'étude de *LinkedIn* de 2015, 61 % des entreprises ont recruté des candidats passifs, ce qui nécessite une exploration efficace pour établir des listes restreintes de candidats.

Les études montrent que le sourcing en ligne, notamment sur LinkedIn, se développe rapidement (Abbas et al., 2021). Cela soulève de nombreuses questions sur la manière dont

les candidats sont triés et sélectionnés parmi la vaste quantité de profils pertinents et accessibles en ligne (D'Silva, 2020; Roulin & Levashina, 2019).

### **Précédemment, dans les études sur le recrutement...**

À notre connaissance, la littérature scientifique ne permet pas encore de savoir si les décisions des recruteurs peuvent être affectées par des sous-optimalités liées à leurs capacités cognitives. Quelques études ont mis en évidence des processus de lecture de CV sous-optimaux (Cole et al., 2007), mais la plupart des recherches sur le recrutement se concentrent davantage sur l'aspect social (comme les écarts de genre ou la discrimination raciale) que sur l'aspect cognitif. Les études en économie, psychologie et sociologie sur le recrutement ont clairement démontré des discriminations systématiques, généralement causées par des stéréotypes, une forme de biais. Les études d'audit et de correspondance ont quantifié la discrimination sur le marché du travail, montrant notamment comment des caractéristiques telles que les noms ou l'origine ethnique influencent les chances d'être rappelé pour un entretien. De nombreux résultats indiquent unanimement que les minorités sont systématiquement désavantagées sur le marché du travail (Eyting, 2022; Jowell & Prescott-Clarke, 1970; Koch et al., 2015; Kroll et al., 2021; Lang & Lehmann, 2012; Neumark, 2018; Pager et al., 2009). Les études futures devraient davantage explorer les solutions pour réduire ces discriminations dans le marché du travail (Bertrand & Duflo, 2017).

Une étude de 2016 a fourni une analyse théorique et expérimentale approfondie du concept de discrimination attentionnelle, qui postule que les décideurs, comme les recruteurs, allouent leur attention de manière endogène et biaisée en fonction des stéréotypes associés à un groupe ethnique ou social et du type de marché (Bartoš et al., 2016). L'article démontre que les candidats soumis à des stéréotypes reçoivent moins d'attention dans les marchés dits de "cherry-picking" – des situations où un filtrage important se produit entre deux étapes. Par exemple, sur le marché du travail, un filtrage significatif a lieu entre la phase de candidature et l'étape des entretiens. En conséquence, les candidats soumis à des stéréotypes, dont les candidatures reçoivent moins d'attention et d'exploration, ont moins de chances de figurer parmi les quelques-uns invités à un entretien. À l'inverse, les individus stéréotypés reçoivent davantage d'attention dans les marchés dits de "lemon-dropping" – des situations où peu de filtrage se produit entre une étape et la suivante. Par exemple, dans le marché de la location immobilière, peu de candidats sont rejetés entre le premier contact et les rendez-vous de visite. Ici, les candidats stéréotypés ont une plus grande probabilité de figurer parmi les quelques-uns empêchés de programmer une visite, car leurs profils reçoivent plus d'attention dans ces situations.

Ces résultats mettent en évidence que la discrimination peut se manifester dès la phase de recherche d'informations, où les recruteurs appliquent des règles de filtrage sous l'influence des stéréotypes. Cette discrimination attentionnelle biaisée, basée sur le type de marché, conditionne la recherche d'informations lors de la sélection d'ensembles, exacerbant ainsi les inégalités dans les situations quotidiennes.

### **Affiner le sujet de recherche**

Notre projet initial visait à étudier les processus réels de recrutement (évaluer les biais cognitifs des recruteurs et leur efficacité dans l'utilisation d'outils informatiques d'évaluation) au sein d'entreprises françaises, par le biais d'un partenariat avec un cabinet de conseil. Cependant, toutes nos propositions ont été refusées pour diverses raisons. Il s'avère que les processus de recrutement sont moins standardisés qu'on pourrait le croire. Parmi les entreprises approchées, très peu disposaient d'archives sur les candidatures passées ou les procédures de recrutement, rendant impossible la reconstitution de tels processus. D'autres entreprises ont décliné par crainte d'accorder l'accès à des données internes sensibles. En conséquence, nous avons décidé de nous concentrer uniquement sur des expériences en laboratoire.

Bien que l'étude des biais sociaux comme les stéréotypes soit bien établie dans la littérature (Bordalo et al., 2016; Cauthen et al., 1971; Ellemers, 2018; Sue & Kitano, 1973), il reste des inconnues significatives sur la manière dont les recruteurs gèrent l'énorme quantité d'informations disponible pour le sourcing. La maîtrise de la recherche d'informations visant à créer une liste restreinte apparaît comme une compétence particulièrement pertinente pour les professionnels du recrutement. Le point de départ de notre paradigme est la question suivante : comment rechercher des informations sur des candidats lorsqu'il est impossible de collecter toutes les informations et qu'il faut établir une liste restreinte ? Nous avons conçu un cadre expérimental pour créer une tâche de recrutement originale capturant la complexité d'une pratique servant de point de départ à de nombreux processus de recrutement.

### **Sélection unique vs sélection d'un ensemble : quelle différence cela fait-il ?**

Avant de présenter en détail notre paradigme dans la section suivante, nous nous concentrons ici sur les différences fondamentales entre le choix d'une seule option et la sélection d'un ensemble d'options. Comme nous allons le voir, la sélection d'un ensemble constitue un problème complexe où la stratégie de recherche peut être contre-intuitive et où les risques d'erreurs et de biais sont accrus par rapport à une sélection unique.

## **Choisir une seule option vs choisir un ensemble d'options**

La prise de décision est un processus omniprésent dans la vie quotidienne et professionnelle, impliquant souvent une recherche d'informations avant de faire un choix. Cela peut être un étudiant sélectionnant les cours qu'il souhaite suivre, un vacancier décidant des activités, quelqu'un choisissant des parfums de glace ou une école sélectionnant ses futurs élèves. Le point commun entre ces exemples est qu'ils décrivent des situations où des individus doivent sélectionner un ensemble d'options. Ces situations diffèrent d'autres, comme une personne choisissant un nouveau téléphone, une famille décidant de sa prochaine destination de vacances ou quelqu'un sélectionnant un film à regarder.

La sélection d'un ensemble d'options implique de multiples erreurs potentielles qui ne surviennent pas dans une sélection unique. Former une équipe de 10 personnes nécessite non seulement de considérer chaque élément individuellement, mais aussi les interdépendances entre eux, ainsi que le besoin de cohésion ou de complémentarité. Par exemple, un étudiant choisissant des cours pour son cursus doit s'assurer que les cours choisis sont compatibles en termes d'emploi du temps. Dans les sports d'équipe, un entraîneur doit construire son équipe tout en veillant à la cohésion entre les joueurs. De même, choisir un ensemble de parfums de glace ne consiste pas seulement à sélectionner ses préférés, mais aussi à s'assurer que les parfums choisis se complètent.

On pourrait supposer que les conséquences d'un mauvais choix sont moins graves dans la sélection d'un ensemble que dans une sélection unique. Malheureusement, ce n'est pas toujours le cas. Faire un mauvais choix parmi 10 pourrait sembler moins critique que de faire un mauvais choix parmi un seul, car les 9 bons éléments restants pourraient théoriquement compenser l'erreur. Cependant, si les bons éléments ne peuvent pas compenser les erreurs, toute la sélection peut échouer. Une équipe dont les membres nuisent à sa cohésion peut perdre un match, un seul mauvais parfum peut ruiner toute la combinaison de glaces, et des cours qui se chevauchent dans un emploi du temps peuvent perturber toute votre organisation. En fin de compte, le risque d'échec est plus élevé dans la sélection d'un ensemble que dans une sélection unique.

## **Trouver une stratégie de recherche d'information adaptée**

Pour faire le meilleur choix, la recherche d'information préalable est cruciale et diffère entre une sélection unique et une sélection d'ensemble. Pour la sélection d'un ensemble, il ne s'agit pas simplement d'identifier la meilleure option unique. Supposons qu'un étudiant doive choisir cinq cours parmi dix options disponibles pour valider son année. Il peut consulter le

syllabus de chaque cours et assister aux présentations des cours. Cependant, il ne peut pas assister à toutes les présentations et doit se fier uniquement au syllabus pour cinq des dix cours. À quelles présentations de cours l'étudiant devrait-il assister ? S'il devait choisir un seul cours, la réponse est simple : après avoir consulté les syllabi, il devrait assister uniquement aux présentations des cours les plus attrayants. Cette approche est intuitive et optimale.

Mais s'il doit choisir cinq cours ? Intuitivement, on pourrait utiliser la même stratégie – assister aux présentations des cinq cours les plus attrayants. Cependant, cette stratégie, bien qu'intuitive, n'est plus optimale. En général, la recherche d'information doit permettre d'identifier le meilleur ensemble d'options. Dans un contexte contraint par le temps, les ressources financières ou la capacité de calcul, la recherche d'informations devrait se concentrer sur les informations les plus décisives pour la sélection d'un ensemble. Dans le cas de la sélection de cinq cours sur dix, l'étudiant bénéficierait le plus en assistant aux présentations des cours pour lesquels il est le plus incertain. Assister aux présentations des cours qu'il est déjà enclin à choisir ou presque certain de rejeter offre peu de valeur pour finaliser sa sélection. En d'autres termes, dans le contexte de la sélection d'un ensemble, se concentrer sur les options les plus préférées devient sous-optimal, ce qui est largement contre-intuitif, et très différent du contexte d'une sélection unique. Les stratégies de recherche d'informations peuvent donc différer profondément entre une sélection unique et une sélection d'ensemble.

### **Le risque de biais et d'heuristiques dans un problème complexe**

Comme le montre la littérature sur la prise de décision et la recherche d'informations, les individus présentent une rationalité limitée (Simon, 1955), caractérisée par des biais cognitifs et l'utilisation d'heuristiques. Les biais cognitifs sont généralement décrits comme des erreurs systématiques et intuitives de raisonnement par rapport à une solution logique ou mathématique. Par exemple, on dit aux participants qu'il y a un groupe de 100 personnes, dont 30 sont ingénieurs. Lorsqu'une personne (nommée Jack) est décrite comme quelqu'un qui aime les puzzles, les participants ont tendance à surestimer la probabilité que Jack soit un des ingénieurs. Cependant, leurs estimations sont exactes lorsqu'aucune description de Jack n'est fournie. L'heuristique associée est "l'heuristique de représentativité", qui consiste à juger la probabilité qu'un événement appartienne à une catégorie en fonction de sa similarité avec le prototype de cette catégorie. Cela conduit à un biais cognitif connu sous le nom de "négligence du taux de base", où les individus ont tendance à ignorer les probabilités préalables lorsqu'ils évaluent la probabilité d'un événement (Kahneman & Tversky, 1973). D'un autre point de vue, les heuristiques sont des stratégies permettant aux individus

d'utiliser efficacement leurs ressources cognitives en simplifiant le problème avant de le résoudre (Gigerenzer et al., 1991).

La taille de la sélection augmente également la complexité computationnelle. Si la complexité d'un choix unique dépasse souvent nos capacités cognitives, elle devient encore plus difficile dans la sélection d'un ensemble. Alors que choisir une option parmi 10 nécessite d'évaluer et de comparer 10 résultats attendus, sélectionner 2 options parmi 10 implique de considérer 45 paires possibles. Pour un étudiant choisissant 5 cours parmi 10, il y a 252 combinaisons possibles. Les défis posés par la sélection d'un ensemble ont été peu explorés, bien qu'il soit évident que notre rationalité limitée peut nous empêcher de résoudre parfaitement de tels problèmes.

### **Question de recherche**

En résumé, la sélection d'un ensemble diffère d'une sélection unique parce que la stratégie de recherche optimale nécessaire est fondamentalement différente et contre-intuitive. La sélection d'un ensemble peut impliquer des risques accrus ou nécessiter de prendre en compte des interdépendances. Enfin, la sélection d'un ensemble est computationnellement plus complexe car elle exige de considérer un nombre de possibilités significativement plus grand et peut être davantage sujette à des stratégies sous-optimales ou biaisées en raison de cette complexité.

La littérature a largement ignoré la manière dont la sélection d'un ensemble met au défi nos capacités cognitives. Cette thèse vise à combler cette lacune et à comprendre les stratégies qu'emploient les individus pour effectuer des sélections d'ensemble. Plus précisément, nous cherchons à comprendre les stratégies d'exploration utilisées pour choisir les informations à examiner et comment les individus effectuent des sélections d'ensemble malgré l'incertitude concernant certaines des options.

## **Méthode Générale : Développement d'un paradigme expérimental pour l'étude de la recherche d'informations dans la sélection d'ensemble**

Dans cette section, nous décrivons le paradigme expérimental que nous avons développé et utilisé tout au long de cette thèse. Plusieurs versions de la tâche ont été employées. Nous présentons ici la version de base qui a servi de fondement à la thèse, en détaillant ses

caractéristiques et le cadre de recherche dans lequel elle a été appliquée. Les versions dérivées de la tâche sont introduites dans les chapitres correspondants.

## Paradigme d'exploration et de sélection d'ensemble

Le paradigme comprend deux étapes : une phase d'exploration suivie d'une phase de sélection. Après avoir reçu une information initiale (un score entier A compris entre 0 et 10) pour chacune des 10 options disponibles, les participants passent à la phase d'exploration. Au cours de cette phase, ils sélectionnent 5 options pour lesquelles ils recevront une information supplémentaire (un autre score entier B compris entre 0 et 10). Après avoir obtenu ce second score pour ces 5 options, les participants entrent dans la phase de sélection. À ce stade, ils sélectionnent les 5 options parmi les 10 initiales qu'ils considèrent comme les meilleures (en se basant sur la somme des scores A et B).

Issu de l'idée initiale d'étudier les décisions de recrutement, les options du paradigme représentent des profils d'individus candidats potentiels à un poste non spécifié. Les scores A et B représentent des compétences distinctes et sont délibérément non corrélés (ce qui est précisé aux participants). La sélection finale lors de la deuxième étape correspond à l'ensemble des candidats recrutés.

## Caractéristiques clés du paradigme

Notre paradigme repose sur quatre caractéristiques essentielles : des actions simultanées, une complexité computationnelle, la présence d'une heuristique intuitive mais sous-optimale, et une heuristique non intuitive mais quasi optimale. Chacune de ces caractéristiques est détaillée ci-dessous.

### **Actions simultanées**

Tout d'abord, le paradigme capture l'exploration et la sélection d'ensemble à travers des actions simultanées à chaque étape. Les participants doivent toujours composer un ensemble avant de passer à l'étape suivante. Cette contrainte les oblige à déterminer un ensemble sans connaître les résultats de leurs actions.

### **Accès limité à l'information**

Ensuite, l'accès à l'information est strictement limité, tout comme la taille de l'ensemble de sélection. Les participants doivent explorer et sélectionner exactement le nombre d'options spécifié (5 explorations et 5 sélections dans la version de base de la tâche). Ils n'ont jamais

accès à 100 % de l'information, mais ils ne peuvent pas non plus choisir d'acquérir moins d'informations. Par conséquent, l'acquisition d'information est gratuite, il n'y a aucune contrainte de temps pour compléter les étapes, et la phase de sélection inclut toujours des scores B inconnus. Cependant, les participants sont libres de choisir pour quelles options ils souhaitent obtenir des informations.

### **Complexité computationnelle**

Le paradigme est également computationnellement complexe. Calculer les gains attendus pour chaque possibilité d'exploration et de sélection afin de déterminer la stratégie optimale est prohibitif, même pour des ordinateurs de laboratoire standard. Par exemple, avec 5 options à choisir parmi 10, il existe 252 combinaisons possibles d'ensembles, contre seulement 10 combinaisons dans un scénario de choix unique. Si les participants doivent explorer 5 options parmi 10 avant de trouver le meilleur ensemble de 5, il existe 252 combinaisons d'exploration, chacune suivie de 252 choix d'ensemble, soit 63 504 trajectoires potentielles. De plus, chaque option comprend deux scores avec 11 valeurs possibles (entiers de 0 à 10). La complexité computationnelle est analysée plus en détail dans le chapitre 1.

### **Heuristiques intuitives et quasi optimales**

Enfin, le paradigme permet aux participants d'utiliser des heuristiques. Plus précisément, il existe une heuristique intuitive qui est clairement sous-optimale. Cette heuristique consiste à explorer les options ayant les scores A les plus élevés, comme mentionné précédemment dans l'exemple d'un étudiant choisissant ses cours. Il existe également une heuristique non intuitive mais quasi optimale, qui consiste à explorer les options situées autour des rangs intermédiaires (selon les scores A), car elles sont les plus susceptibles d'entrer ou de sortir de l'ensemble de sélection final. Par exemple, dans le cas d'un étudiant sélectionnant 5 cours, la meilleure heuristique consisterait à assister aux présentations des cours situés près de la 5e position après avoir examiné toutes les descriptions des cours. Dans l'expérience, il est donc plus efficace d'explorer les scores B des candidats proches du 5e rang selon leurs scores A. Cette heuristique peut être ajustée pour s'adapter à la taille des ensembles d'exploration et de sélection (voir chapitres 1 et 4 pour différentes versions de la tâche). Ces heuristiques peuvent être appliquées à chaque essai, sans nécessiter d'apprentissage d'un essai à l'autre, car toutes les données du problème sont connues.

### **Cadre de recherche du paradigme**

Ce paradigme étant inédit et traitant d'une question originale dans la littérature, la thèse se concentre sur quelques questions de recherche fondamentales liées à la sélection d'ensemble.

Nous privilégions l'examen de l'écart entre la solution optimale et les stratégies des participants. Nous décrivons les comportements des participants en utilisant des heuristiques simples pour tirer des conclusions initiales. Les phénomènes observés sont comparés à plusieurs biais cognitifs bien documentés dans la littérature. Enfin, nous cherchons à comprendre comment les individus perçoivent leur propre stratégie pour résoudre le paradigme avant de tester des méthodes potentielles de remédiation.

## Méthodologie expérimentale (du Chapitre 1)

### Stimuli et procédure

La tâche principale des participants était présentée comme une tâche de recrutement : à chaque essai, ils étaient confrontés à 10 profils et devaient identifier les 5 meilleurs parmi ces 10 profils. Chaque profil était composé de 2 scores, notés A et B. Tous les scores étaient tirés indépendamment d'une distribution uniforme entre 0 et 10, et cette information était connue des participants. Au début de chaque essai, tous les scores A étaient révélés, mais les scores B restaient cachés. Lors de la phase d'exploration, les participants devaient choisir 5 profils pour lesquels la valeur du score B serait révélée. Ensuite, lors de la phase de sélection, les participants devaient sélectionner les 5 meilleurs profils, définis comme ceux ayant les moyennes les plus élevées des scores A et B. Il était précisé aux participants que les scores A et B avaient une importance égale, même pour les scores non révélés.

Après les instructions, les participants réalisaient un essai d'entraînement, au cours duquel tous les scores étaient révélés après la phase de sélection, afin de s'assurer qu'ils comprenaient bien la tâche et l'importance des informations non révélées.

Après chaque essai, les participants recevaient un retour sur la performance de leurs choix, c'est-à-dire le nombre de profils dans leur sélection qui appartenait à l'ensemble des meilleurs profils. Il est à noter qu'en cas d'égalité, l'ensemble des meilleurs profils pouvait dépasser 5.

Chaque participant effectuait un total de 25 essais, en 25 à 30 minutes (instructions incluses). Les participants pouvaient recevoir un bonus lié à leur performance : un essai était sélectionné aléatoirement, et ils recevaient 1 euro pour chaque bonne réponse dans cet essai.

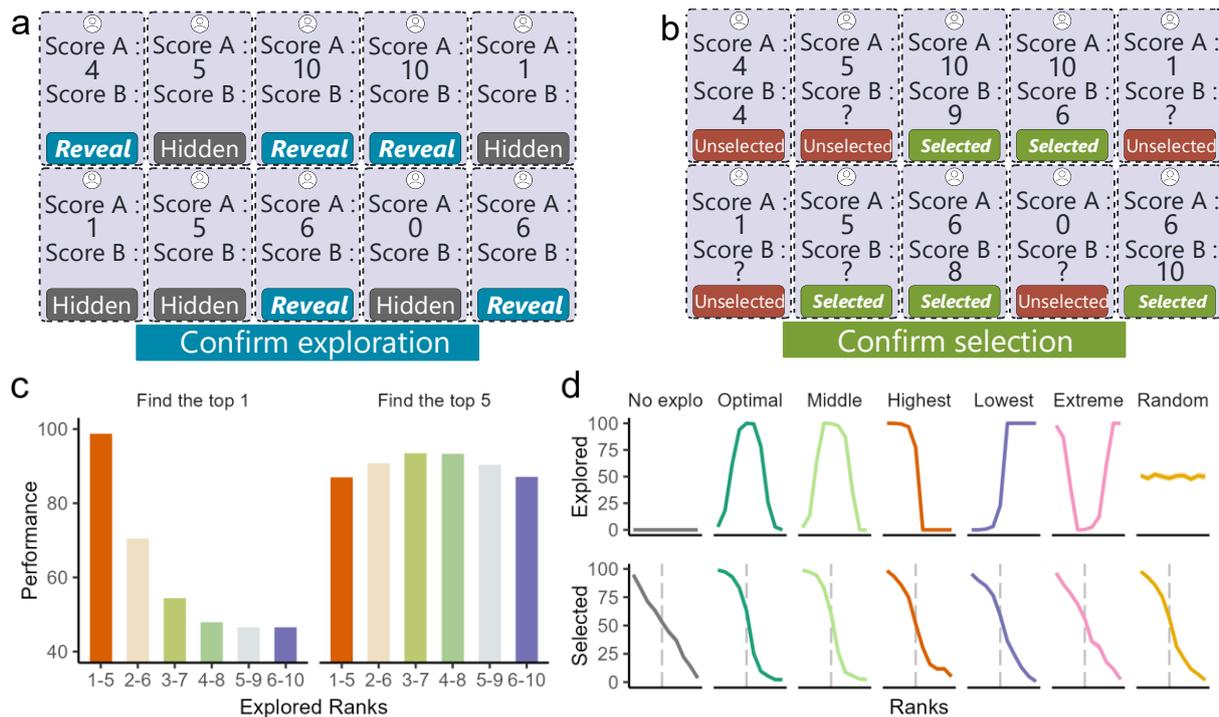


Fig 1.

Présentation de la tâche principale et conséquences de la sélection d'ensemble selon la stratégie utilisée. La tâche principale comporte deux étapes pendant l'expérience : (a) Phase d'exploration : les participants consultent les scores A des 10 options disponibles et indiquent 5 options pour lesquelles ils souhaitent révéler les scores B, initialement cachés. (b) Phase de sélection : les participants accèdent aux scores B qu'ils ont demandés et doivent sélectionner les 5 meilleures options parmi les 10, sur la base de la somme des scores A et B. Comment explorer lorsque l'objectif est de sélectionner la meilleure option unique par rapport aux 5 meilleures options ? Nous avons calculé (c) la performance d'un agent virtuel explorant les 5 options ayant les scores A les plus élevés, ou les scores de rangs intermédiaires (par exemple, les rangs 2 à 6), et effectuant une sélection idéale basée sur ces informations. Lorsque l'objectif est de sélectionner la meilleure option unique, la stratégie optimale consiste à explorer les rangs 1 à 5. En revanche, lorsqu'il s'agit de sélectionner les 5 meilleures options, il devient optimal d'explorer les rangs intermédiaires. Enfin, nous avons illustré (d) diverses stratégies d'exploration. Les graphiques montrent la proportion de profils explorés et sélectionnés en fonction de leur rang selon le score A (phase d'exploration, graphiques supérieurs) ou le score total (phase de sélection, graphiques inférieurs). Les stratégies d'exploration incluent : No-explo : l'agent ne reçoit aucune information pendant la phase d'exploration. Optimal : l'agent effectue une exploration optimale à chaque essai. Les stratégies axées sur les 5 meilleures options selon le score A (*Highest*), les 5 moins bonnes (*Lowest*), les options de rangs intermédiaires (*Middle*), les options aux rangs extrêmes (*Extreme*) ou une exploration aléatoire de 5 options (*Random*). Chaque stratégie d'exploration implique une sélection optimale basée sur les informations explorées.

## Modèle et mesures

### Calcul de la performance

La performance d'une sélection est définie comme la proportion d'options sélectionnées qui vérifient l'objectif. La performance normalisée a été calculée comme la performance du participant moins la performance aléatoire, divisée par l'écart entre la performance optimale et la performance aléatoire.

## Stratégies d'exploration

Pour caractériser le comportement des participants pendant la phase d'exploration, nous avons comparé leurs choix d'exploration à plusieurs heuristiques :

- Highest : exploration des options ayant les meilleurs rangs.
- Middle : exploration des options de rangs intermédiaires.
- Lowest : exploration des options ayant les rangs les plus faibles.
- Extreme : exploration des options aux rangs les plus élevés et les plus faibles.

À noter que l'heuristique Middle inclut un choix aléatoire entre l'option classée 3e et celle classée 8e, en plus d'explorer toutes les options classées entre la 4e et la 7e position. De même, pour l'heuristique Extreme, un choix est fait entre l'option classée 3e et celle classée 8e, en plus des explorations des options aux deux premiers et deux derniers rangs.

Nous avons également pris en compte une stratégie d'exploration aléatoire (Random), une stratégie sans exploration (No Expl), et estimé la stratégie d'exploration optimale (Optimal), comme détaillé ci-dessous et dans les matériaux supplémentaires.

Sélections optimales conditionnelles à l'exploration

Lors de l'évaluation des performances des différentes stratégies d'exploration, nous avons pris en compte leur impact sur les sélections finales. En effet, la meilleure approche de sélection dépend des informations obtenues pendant la phase d'exploration. Pour chaque stratégie d'exploration, nous avons donc considéré la sélection qui maximiserait la performance attendue.

## Exploration optimale

L'exploration optimale est définie comme celle menant à une performance attendue maximale lorsqu'elle est suivie d'un processus de sélection optimal. Malgré la simplicité apparente du paradigme, nous n'avons pu estimer la stratégie d'exploration optimale qu'à l'aide d'une méthode d'échantillonnage (voir détails dans les matériaux supplémentaires). Comme illustré dans la figure 1d, la stratégie optimale dans le cadre de l'expérience 1 est proche de l'heuristique Middle. Cependant, elle n'est pas identique, car elle prend en compte à la fois les rangs des options et leurs valeurs numériques. Par exemple, si l'écart entre les scores A des options classées 6e et 8e est important, mais que celui entre les options classées 3e et 5e est faible, l'exploration optimale inclura l'option classée 3e mais pas la 8e. En revanche, l'heuristique Middle, qui ne considère que les rangs, choisirait aléatoirement entre les deux options.

Dans la figure 1a, la stratégie optimale consiste à explorer les scores des options classées 4, 5, 5, 6 et 6, tandis que la stratégie Middle explore également les scores des options 4, 5, 5 et 6, mais choisit au hasard entre la deuxième option ayant un score de 6 et une autre ayant un score de 1.

La stratégie optimale décrite ici surpasserait systématiquement toute autre heuristique dans les cas où les scores A et B sont corrélés. Par exemple, la stratégie Highest n'est pas optimale lorsque les scores sont corrélés.

### **Calcul des biais et des erreurs**

Pendant la phase d'exploration, le biais en faveur des options les mieux classées est défini comme la différence entre la proportion d'options ayant les meilleurs scores A explorées par les participants et la même proportion évaluée pour la stratégie d'exploration optimale. Une valeur positive indique que les participants explorent davantage les options ayant les meilleurs scores A que ce qui est prescrit par la solution optimale.

Pendant la phase de sélection, un biais en faveur des options explorées est défini comme la différence entre la proportion d'options explorées choisies par les participants et la même proportion évaluée pour la solution optimale (en tenant compte des explorations des participants). Une valeur positive indique que les participants montrent un biais en faveur des options explorées pendant la phase de sélection.

Nous avons également évalué les "erreurs mathématiques" commises par les participants lors de leur sélection, séparément pour les options explorées et non explorées. Pour ce faire, nous avons vérifié si, lorsqu'ils sélectionnaient des options parmi celles explorées (respectivement non explorées), ils choisissaient effectivement celles ayant les scores totaux les plus élevés (respectivement le score A le plus élevé). Nous avons additionné les erreurs commises pour les options explorées et non explorées, puis divisé cette somme par le nombre total d'options sélectionnées, afin d'exprimer ces erreurs en pourcentage.

## **Chapitre 1 : Biais d'exploration en faveur des options favorites dans la sélection d'ensemble**

Ce chapitre examine comment les individus abordent la recherche d'informations lorsqu'ils doivent sélectionner un ensemble d'options. Nous comparons la stratégie d'exploration des

participants à la stratégie optimale ainsi qu'à plusieurs stratégies simples. L'étude identifie deux biais majeurs : un biais d'exploration, où les participants examinent de manière disproportionnée les options perçues comme favorites pour la sélection finale, et un biais de sélection, où les options explorées sont privilégiées même lorsqu'elles ne sont pas optimales. À travers trois expériences, le chapitre démontre comment ces biais nuisent à la performance.

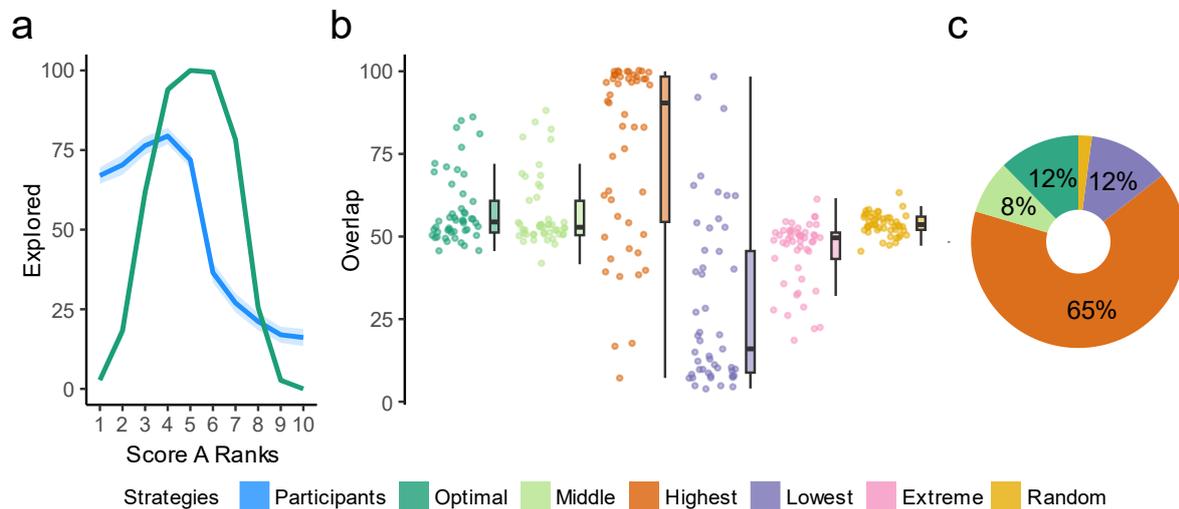


Fig. 2.

Comparaison de l'exploration des participants avec la stratégie optimale et diverses heuristiques dans l'Expérience 1. Les graphiques représentent : (a) les proportions de profils explorés pour chaque rang, pour les participants et pour l'agent optimal, (b) les proportions de recoupement entre l'exploration des participants et celle optimale ou les heuristiques, et (c) les proportions des stratégies les plus proches, parmi les participants.

## Chapitre 2 : Facteurs cognitifs influençant les biais dans la recherche d'informations

Le chapitre 2 explore les bases cognitives des biais d'exploration et de sélection observés dans le premier chapitre. Afin de comprendre la variabilité interindividuelle des biais relevés dans notre paradigme, nous utilisons un ensemble de tâches servant de points de comparaison.

Ce chapitre évalue les relations entre ces biais et des mesures traditionnelles des capacités cognitives (Cognitive Reflection Test, Berlin Numeracy Test) ainsi que des biais cognitifs (par exemple, biais de confirmation, biais d'ancrage, effet de cadrage, erreur de conjonction, négligence du taux de base, aversion au risque). Bien que les scores obtenus dans ces tâches soient cohérents avec la littérature, les résultats révèlent des associations limitées entre ces biais et les biais d'exploration et de sélection dans notre paradigme.

Étant donné que notre tâche s'inspire du sourcing et est présentée comme une expérience de recrutement, nous nous demandons si les experts dans ce domaine manifestent les mêmes biais que les non-experts. Nous recrutons un panel de recruteurs professionnels via LinkedIn pour réaliser une version abrégée de l'expérience. Nous comparons leurs réponses à celles d'une population non experte. Les résultats montrent que l'expertise ne conduit pas nécessairement à l'adoption de stratégies d'exploration plus optimales.

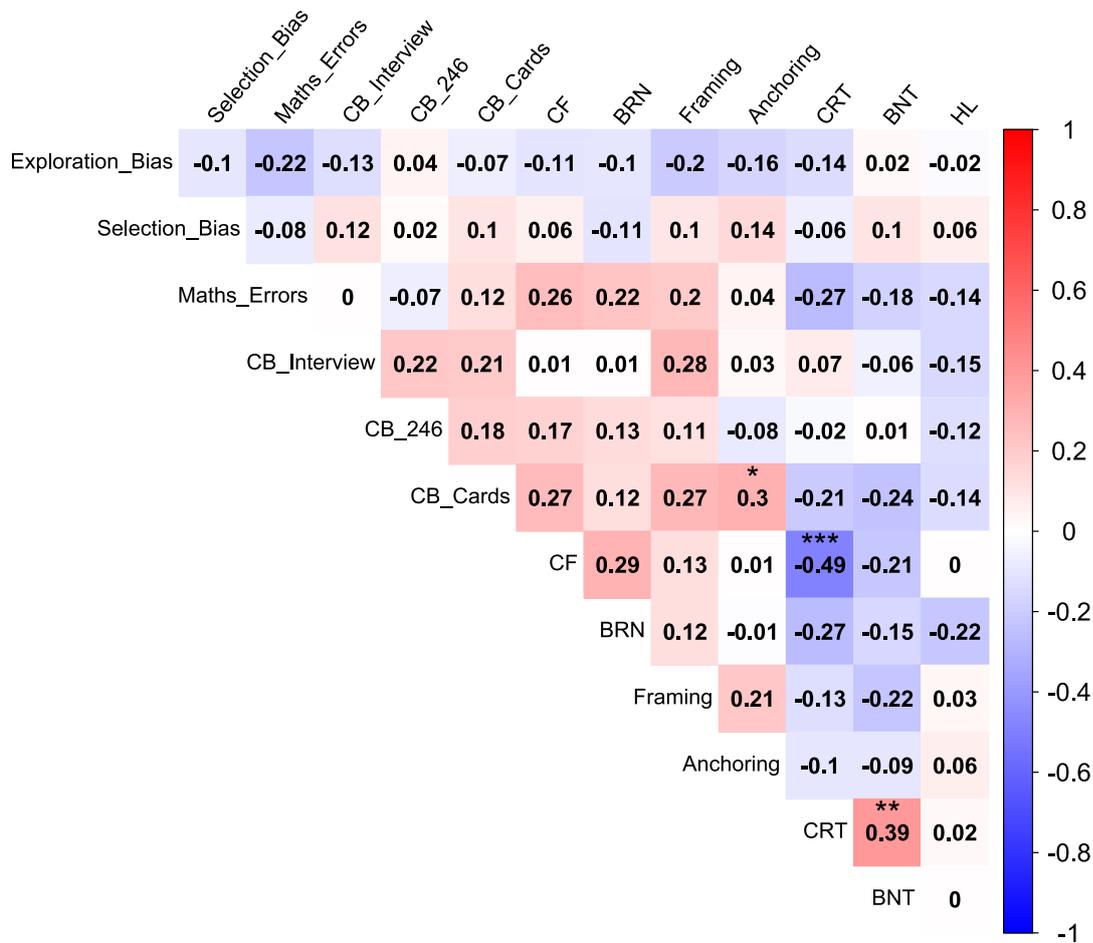


Fig. 3.

Matrice de corrélation entre les tâches utilisées (CRT, BNT, CB-246, CB-Interview, CB-Cards, HL, BRN, CF, Ancrage, Cadrage), les biais d'exploration, les biais de sélection et les erreurs mathématiques dans la tâche de sélection d'ensemble. Niveaux de signification après correction selon la méthode de Benjamini-Hochberg : < .05 \* < .01 \*\* < .001 \*\*\*

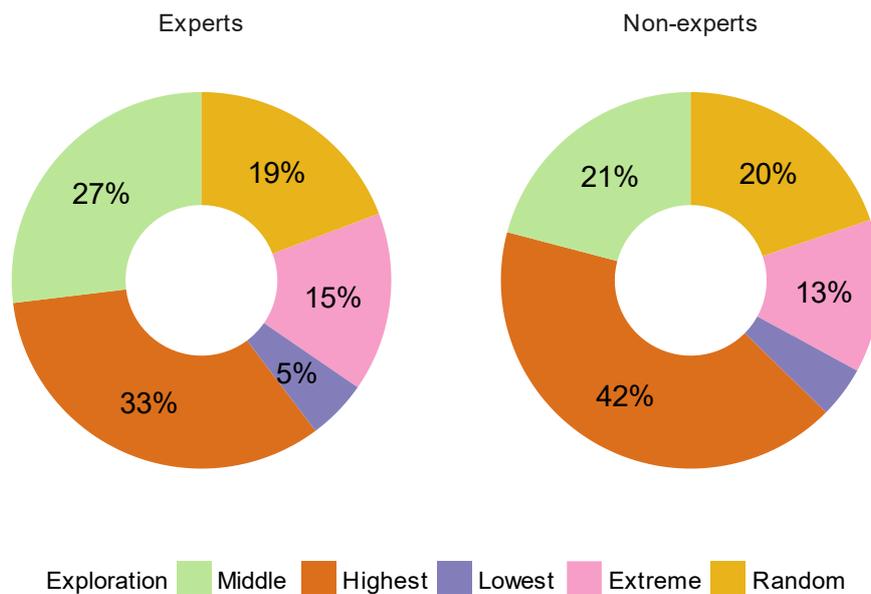


Fig. 4.

Graphiques représentant les choix d'heuristiques d'exploration en fonction de l'expertise en recrutement dans l'Expérience 5. Les choix possibles étaient : Highest (explorer les candidats ayant les scores A les plus élevés), Lowest (explorer les scores les plus bas), Middle (explorer les scores intermédiaires), Extreme (explorer à la fois les scores les plus élevés et les plus bas) ou Random (explorer de manière aléatoire).

## Chapitre 3 : Jugements métacognitifs sur les stratégies d'exploration pour la sélection d'ensemble

À partir des résultats des chapitres 1 et 2, il apparaît que la stratégie des individus est stable d'un essai à l'autre et qu'ils peuvent identifier l'heuristique qui décrit le mieux leur stratégie. Nous nous demandons donc s'ils sont capables d'évaluer l'efficacité de leur stratégie.

Le chapitre 3 examine comment les individus perçoivent et évaluent leurs propres stratégies de recherche d'informations. Il met en évidence une tendance des participants à surestimer l'efficacité de leurs approches préférées, même lorsque celles-ci sont manifestement sous-optimales. Cependant, leur confiance augmente lorsqu'ils appliquent avec succès la stratégie qu'ils ont choisie. Ces résultats soulignent le rôle de la métacognition dans la formation des comportements d'exploration et dans la confiance dans les décisions.

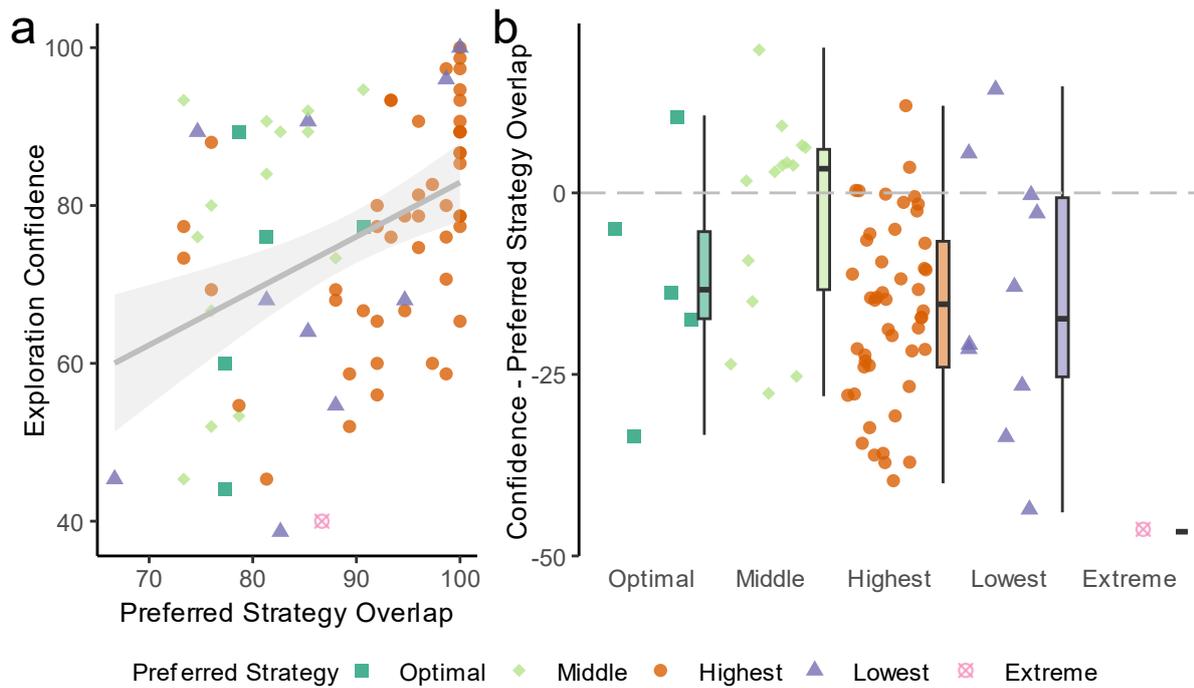


Fig. 5.

Le panneau (a) illustre l'estimation moyenne, par chaque participant, de la qualité de leur exploration en fonction du pourcentage moyen de recouvrement avec leur stratégie préférée. Le panneau (b) montre la différence (en points de pourcentage) entre leur confiance moyenne et le recouvrement moyen avec leur stratégie préférée, selon leur stratégie préférée.

## Chapitre 4 : Apprentissage et adaptation des stratégies optimales

Le chapitre 3 ayant montré que les individus surestiment leur efficacité, la question se pose de savoir comment les aider au mieux. Parmi les moyens possibles de remédiation, une approche éducative semble être l'option la plus efficace lorsqu'elle peut être mise en œuvre. Les incitations financières peuvent également aider à motiver les participants à rester concentrés et à éviter des erreurs de calcul.

Le chapitre 4 explore la capacité d'apprentissage et d'adaptation des individus dans les deux phases du paradigme. Il montre que les individus peuvent améliorer leurs stratégies d'exploration et de sélection lorsqu'ils reçoivent du contenu éducatif accompagné d'un entraînement avec retour d'information. Les participants formés ont tendance à généraliser leurs connaissances à différentes versions des tâches. Ce chapitre examine également l'impact des incitations financières et ne trouve aucun effet significatif sur les stratégies ou les performances.

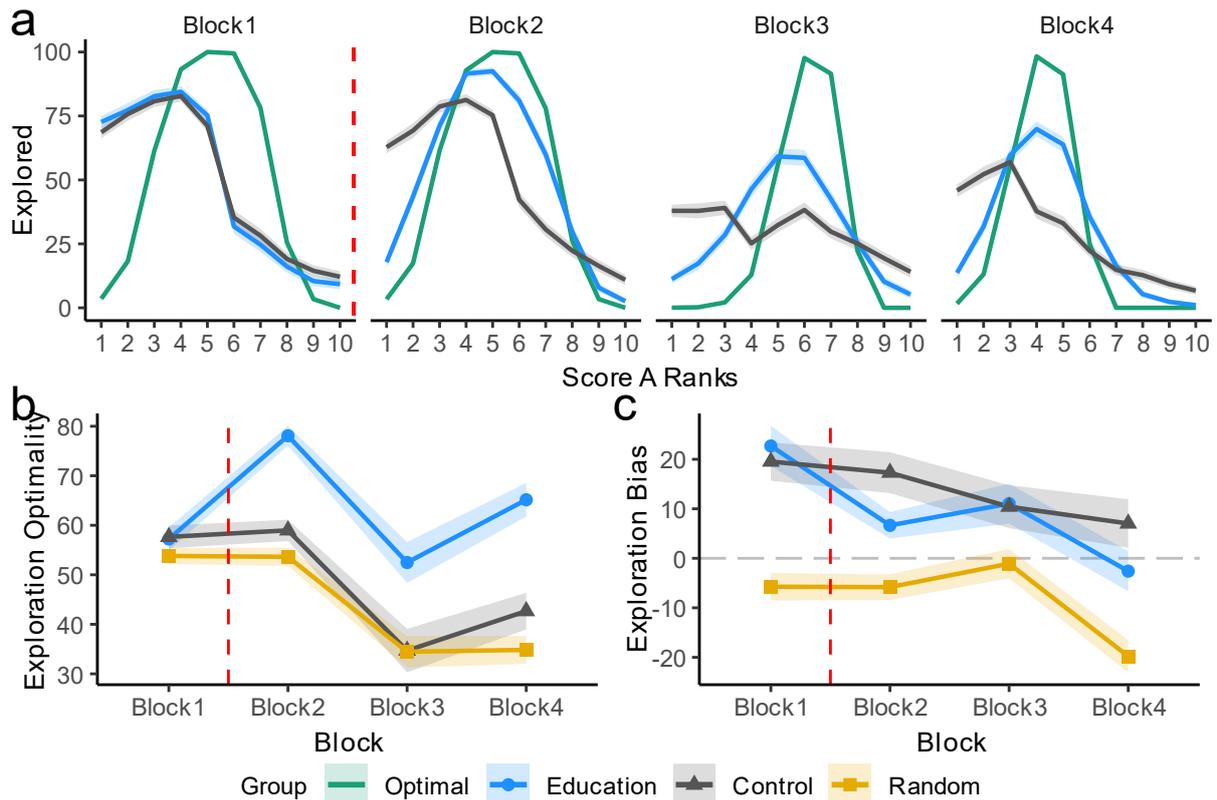


Fig. 6.

Le panneau (a) représente les options explorées par les participants et par la stratégie d'exploration optimale en fonction des rangs des scores A des options, pour chaque bloc. Le panneau (b) montre l'optimalité de l'exploration réalisée par les participants du groupe de remédiation (en bleu) et du groupe contrôle (en gris). L'optimalité d'une exploration aléatoire est représentée (en jaune). Le panneau (c) illustre le biais d'exploration envers les options favorites pour le groupe de remédiation (en bleu) et le groupe contrôle (en gris) à travers les blocs. Le biais d'exploration d'une exploration aléatoire est représenté (en jaune). Les lignes rouges en pointillés représentent le moment de l'intervention.

## Discussion

### Contributions spécifiques

Cette thèse apporte plusieurs contributions distinctes à la littérature sur la recherche d'informations dans la prise de décision, comme détaillé ci-dessous.

Premièrement, nous introduisons une nouvelle question théorique en examinant la sélection d'un ensemble d'options. Cette problématique n'avait pas été abordée dans la littérature, malgré sa pertinence écologique. Comprendre comment les individus résolvent ce type de problème est crucial, étant donné sa complexité computationnelle et les stratégies contre-intuitives qu'il nécessite par rapport à la sélection d'une seule option.

Deuxièmement, nous avons développé un paradigme expérimental spécifique pour répondre à notre question de recherche, plutôt que d'adapter un cadre existant. L'objectif était de simuler une tâche de sourcing en laboratoire, une activité courante pour les recruteurs. Ce paradigme a été conçu dans une approche économique, visant à caractériser la structure de l'information disponible pour les participants et à mesurer l'écart entre les comportements observés et optimaux. Comme démontré tout au long de cette thèse, le paradigme est flexible (par exemple, la taille de l'ensemble d'exploration, la taille de l'ensemble de sélection, le cadrage, les versions des questionnaires). Nous espérons qu'il servira de cadre de référence pour décrire et étudier la recherche d'informations dans des scénarios de sélection d'ensemble.

Troisièmement, cette thèse adopte une approche interdisciplinaire. Le travail s'appuie principalement sur la littérature en psychologie cognitive et en économie expérimentale. La conception de la tâche suit des méthodologies typiques de l'économie expérimentale, telles que l'incitation des choix (mais aussi des croyances rapportées, comme la confiance dans le chapitre 3), l'utilisation de gains probabilistes et de récompenses basées sur des essais choisis aléatoirement. Nous avons également utilisé des outils conceptuels et expérimentaux issus de la psychologie cognitive, comme la classification des biais cognitifs (Ceschi et al., 2019), le Cognitive Reflection Test (Brañas-Garza et al., 2019 for a review; Frederick, 2005) et le Berlin Numeracy Test (Cokely et al., 2012), largement utilisés dans les études de prise de décision dans divers domaines. À la croisée de l'économie expérimentale et des sciences cognitives se trouve également l'utilisation de modèles computationnels et la comparaison des performances des participants avec celles d'agents optimaux. Enfin, dans le chapitre 4, nous avons simultanément testé des méthodes de remédiation cognitive (interventions éducatives) et économique (incitations financières), ce qui nous a permis d'étudier leurs effets indépendants et combinés.

## Limites méthodologiques

Cette thèse repose principalement sur des expériences en laboratoire, qui, bien qu'ineestimables pour une investigation contrôlée, peuvent limiter la validité écologique des résultats. Les tâches conçues, comme le paradigme en deux étapes (exploration et sélection), simplifient les complexités du monde réel pour se concentrer sur des mécanismes cognitifs spécifiques. Ces simplifications incluent, par exemple, l'absence de corrélation entre les scores A et B, le fait que les participants reçoivent les informations sur les scores B seulement après avoir choisi toutes les options explorées (plutôt que d'explorer une option à la fois) et l'utilisation d'un nombre fixe d'options à explorer. Ces simplifications peuvent ne pas

capturer pleinement la dynamique de la prise de décision dans des contextes pratiques, tels que le recrutement ou les processus de sélection académique.

De plus, les échantillons de participants étaient principalement composés d'adultes français provenant d'un unique pool expérimental à Paris (à l'exception de l'échantillon d'experts). Toutes les expériences ont été menées en ligne. Cela limite la généralisation des résultats à des populations et contextes divers. Bien que nos résultats n'aient montré aucun effet significatif lié à l'expertise, à l'âge, au niveau d'éducation ou au genre, nous reconnaissons que de futures recherches devraient viser à reproduire ces résultats avec des populations plus variées et dans des environnements réels pour évaluer la robustesse des biais et heuristiques identifiés.

## Autres Limitations

### Différences inter-individuelles

Nos résultats mettent en évidence des différences individuelles qui pourraient être explorées dans de futures recherches. Nous avons observé que certains participants ne manifestaient ni biais d'exploration ni biais de sélection. De manière intéressante, certains participants présentaient même des biais opposés : un biais d'exploration envers des options externes et un biais de sélection en faveur des options non explorées. Les différences interindividuelles sont également évidentes dans l'évaluation par les participants de la qualité de leur exploration : certains surestiment leur performance, tandis que d'autres la sous-estiment (chapitre 3). Par ailleurs, une partie des participants n'a pas réussi à améliorer sa stratégie d'exploration malgré l'intervention éducative, tandis que d'autres n'ont nécessité aucune intervention et ont démontré une stratégie efficace pour la tâche initiale ainsi que ses variations (chapitre 4).

L'intégration de mesures psychologiques et cognitives, telles que des évaluations de la personnalité ou des tests neuropsychologiques, pourrait aider à examiner comment les traits de personnalité influencent les comportements d'exploration et de sélection. Des études récentes montrent que les différences interindividuelles dans l'expression des biais cognitifs pourraient être liées aux traits de personnalité (Ahmad, 2020; Kumar et al., 2023; Singh et al., 2023). Par exemple, l'extraversion est fortement liée aux biais de surconfiance et de confirmation (Ahmad, 2020; Kumar et al., 2023). En revanche, le neuroticisme conduit à des stratégies plus prudentes face au risque (Singh et al., 2023). Explorer les liens avec les traits de personnalité pourrait aider à mieux décrire les comportements dans notre paradigme.

## Valeur non instrumentale de l'information

La stratégie optimale suppose un cadre utilitaire spécifique, qui peut ne pas s'aligner avec les objectifs individuels des participants. Par exemple, certains participants pourraient privilégier la valeur non instrumentale de l'information — la satisfaction intrinsèque ou la confiance acquise en obtenant de l'information — plutôt que la maximisation des gains attendus (Eliaz & Schotter, 2010; Kobayashi et al., 2019). Dans de tels cas, les participants pourraient se concentrer sur l'amélioration de leur confiance dans leurs pré-décisions (c'est-à-dire les choix qu'ils anticipent de faire avant l'exploration) plutôt que de poursuivre des objectifs purement instrumentaux.

Pour examiner si le biais d'exploration pourrait être appliqué à une tâche dans laquelle l'information n'a pas de valeur instrumentale, nous pourrions concevoir une tâche où la phase d'exploration aurait lieu après la phase de sélection. Par exemple, les participants pourraient effectuer leur sélection en se basant uniquement sur les scores A, suivie d'une phase d'exploration pour obtenir des informations sur les scores B de cinq options. À la fin de l'essai, nous leur fournirions un retour sur la qualité de leur sélection, garantissant que la phase d'exploration n'est pas essentielle pour déterminer la qualité de leur décision.

## Perspectives sur les impacts dans la vie réelle

Sur la base des résultats de cette thèse, nous pouvons envisager les impacts potentiels des enseignements sur des situations réelles, telles que le recrutement ou les algorithmes de recommandation. Nous revisitons le domaine du recrutement, qui a servi de contexte à cette recherche, et discutons des implications potentielles dans le domaine des algorithmes de recommandation. Nous pensons que les biais dans la sélection d'ensemble n'ont pas encore été reconnus comme des problèmes significatifs, malgré leurs effets potentiellement importants.

Nous décrivons l'impact que ces biais peuvent avoir lorsqu'ils interagissent avec d'autres types de biais (stéréotypes, biais algorithmiques) et présentons des stratégies potentielles pour y remédier. Ces implications soulignent la pertinence pratique de cette thèse face aux défis du monde réel, en proposant des stratégies concrètes pour améliorer la qualité des décisions dans des environnements complexes. En abordant les biais inhérents à la sélection d'ensemble, les organisations et les professionnels peuvent promouvoir des systèmes de prise de décision plus justes et plus efficaces.

## Sélection du personnel

### Tri des CV

De nombreuses études mettent en évidence la discrimination dans le tri des CV, soulignant son rôle dans la réduction des chances des candidats d'obtenir un entretien d'embauche et dans la création d'inégalités d'accès à l'emploi (Bertrand & Duflo, 2017; Eytting, 2022; Kroll et al., 2021; Neumark, 2018; Pager et al., 2009). Cette première étape du processus de recrutement consiste généralement à sélectionner des ensembles de candidats potentiels parmi un très grand pool. À cet égard, elle correspond au paradigme présenté dans cette thèse.

Des solutions encourageantes ont été proposées pour réduire la discrimination à cette étape. Certaines études ont montré les effets positifs des interventions visant à aider les recruteurs à réduire leurs stéréotypes, bien que les effets à long terme soient mitigés (Devine et al., 2012; Forscher et al., 2017). Une solution testée est l'introduction de CV anonymes. Là encore, les effets sont mitigés : des effets positifs sont observés, en particulier pour les femmes (Åslund & Skans, 2012), mais des effets négatifs inattendus apparaissent pour les personnes ayant des parcours moins conventionnels, en partie à cause de discriminations passées (Behaghel et al., 2015). Globalement, cette première étape de tri, censée aboutir à une sélection de tous les candidats, n'a pas encore été dé-biaisée et engendre d'importantes inégalités dans l'accès à l'emploi (Petit et al., 2013).

### Interactions avec le biais d'exploration et conséquences

Les biais, tels que les stéréotypes dans les processus de recrutement (Bosak & Sczesny, 2011; Koch et al., 2015; Sczesny et al., 2004) peuvent interagir avec les biais d'exploration. Les stéréotypes influencent le processus de recrutement avant même que les candidats ne soient évalués, agissant comme un filtre initial pour les recruteurs. Dans notre paradigme, les stéréotypes fonctionnent de manière similaire au "score A". Au lieu de recevoir un score initial objectif, les recruteurs se fient à une évaluation subjective biaisée. En conséquence, les stéréotypes déterminent quels candidats sont considérés comme favoris pour la sélection.

Lorsque les stéréotypes orientent l'identification des favoris, les individus ciblés par ces biais ont moins de chances de voir leur candidature explorée (Bartoš et al., 2016). Cela aggrave le problème, car leurs CV sont également évalués moins favorablement lorsqu'ils sont examinés. Par conséquent, le préjudice est double : moins d'opportunités d'être considéré dans l'ensemble de sélection et moins de chances d'être évalué positivement s'ils le sont.

Globalement, cela explique pourquoi les personnes victimes de stéréotypes sont rejetées, même à des étapes où un large éventail de candidats est encore en course.

Des stratégies peuvent être conçues pour atténuer ces effets. Par exemple, la stratégie d'exploration optimale décrite dans cette thèse peut ne pas être réalisable dans certains contextes. Comme discuté dans le chapitre 1, les procédures de sélection des étudiants pour les programmes académiques excluent souvent un candidat ayant un dossier académique moyen mais une excellente performance à l'entretien, au profit d'un autre candidat ayant un dossier légèrement meilleur mais qui n'a pas été présélectionné pour un entretien.

D'autres solutions pourraient traiter les effets combinés de ces pratiques. En France, les quotas ne peuvent légalement pas être appliqués à des critères tels que l'origine ethnique ou l'âge, contrairement au genre ou au handicap. Une solution plus générale pourrait consister en un droit à une seconde chance. À chaque étape du processus de recrutement, des candidats exclus aléatoirement pourraient être réintroduits dans les étapes suivantes. Cette approche permettrait de corriger les erreurs des recruteurs et offrirait aux candidats une autre opportunité d'être considérés. Une telle solution compléterait d'autres efforts pour améliorer les processus de recrutement, tels que la formation des recruteurs, l'utilisation d'entretiens structurés ou de tests de validité prédictive.

## Algorithmes de recommandation

Les algorithmes de recommandation sont omniprésents dans notre quotidien, intégrés dans presque tous les outils numériques. Leur principal objectif est de nous fournir du contenu pertinent pour nous aider à prendre les bonnes décisions. Sur les réseaux sociaux comme Facebook, Instagram, X et TikTok, ou sur les services de streaming comme YouTube, Netflix et Spotify, ils façonnent le contenu que nous voyons, en l'adaptant à nos préférences. Sur les sites de commerce en ligne comme Amazon ou Shein, ils suggèrent des produits que nous sommes les plus susceptibles d'acheter. Les sites d'actualités les utilisent pour recommander des articles liés à celui que nous venons de lire. Même dans le domaine scientifique, les chercheurs s'appuient fortement sur les recommandations de l'algorithme de Google Scholar après des recherches par mots-clés.

## Biais algorithmiques

Les biais algorithmiques, définis comme des distorsions systématiques dans les résultats produits par les systèmes automatisés, influencent la manière dont les utilisateurs perçoivent et interagissent avec le contenu. Par exemple, le biais de popularité dans les

systèmes de recommandation amplifie la visibilité des éléments déjà populaires au détriment des moins visibles, créant un effet "les riches deviennent plus riches et les pauvres plus pauvres" (Abdollahpouri, 2019). Un autre exemple est fourni par Cho et al. (2020), qui montrent comment les algorithmes de YouTube exacerbent la polarisation politique en recommandant des vidéos alignées sur les préférences idéologiques de l'utilisateur. Ces biais ne se limitent pas aux données, mais sont également intégrés dans les modèles (Kordzadeh & Ghasemaghaei, 2022).

While several solutions have been proposed to mitigate algorithmic biases, their effectiveness remains mixed depending on the context and implementation. For instance, techniques such as reweighting data or integrating fairness metrics into algorithms, as suggested by Chen et al. (2023), have shown potential in reducing popularity bias and improving exposure diversity. However, these approaches often face trade-offs, such as reduced overall accuracy or increased computational complexity (Abdollahpouri & Mansoury, 2020). User-driven solutions, like the interactive tools tested by Millecamp et al. (2018) on Spotify, empower individuals to control recommendations, yet their success depends heavily on user engagement and understanding. Transparency and audit practices have been effective in identifying systemic issues (Kordzadeh & Ghasemaghaei, 2022), but implementing these measures at scale remains a challenge, highlighting the need for more robust, context-specific interventions.

### **Interactions avec le biais d'exploration et conséquences**

Bien que les interactions homme-machine soient fondamentales dans l'utilisation des algorithmes de recommandation, il est frappant de constater que leur conception néglige souvent la complexité du comportement humain. Le biais de popularité illustre clairement comment les préférences humaines pour les favoris sont reflétées dans les systèmes de recommandation. En intégrant un biais humain dans un algorithme de recommandation, ce biais est effectivement amplifié. En conséquence, les individus sont exposés uniquement à du contenu populaire ou favorisé, limitant leur exploration à ces options avant de sélectionner finalement du contenu qu'ils connaissent déjà bien. Cette polarisation vers du contenu populaire découle non seulement de l'algorithme lui-même, mais aussi de l'interaction entre un individu biaisé et un algorithme tout aussi biaisé.

Les algorithmes de recommandation peuvent être conçus pour réduire activement nos biais. Les propositions actuelles visant à corriger les biais algorithmiques cherchent souvent à éliminer complètement ces biais, en aspirant à une neutralité qui améliorerait l'efficacité des systèmes de recommandation. Cependant, il est possible d'aller plus loin. En adoptant une

approche interdisciplinaire combinant les sciences de l'information et les sciences comportementales, les algorithmes pourraient être intentionnellement conçus pour contrer nos biais. Par exemple, les algorithmes pourraient recommander volontairement des "outsiders" ou du contenu inconnu. Une telle approche contre-intuitive nécessiterait de présenter ces recommandations de manière transparente (pour éviter de tromper les utilisateurs) mais aussi attrayante. En pratique, les algorithmes pourraient encadrer le contenu comme des opportunités pour devenir cinéphile en recherchant une nouvelle série sur Netflix, suggérer des articles pour "penser autrement" en parcourant des articles de presse, ou promouvoir l'idée de devenir "fan précoce" d'un artiste émergent grâce à des recommandations inattendues sur Spotify.

# Bibliographie

1. Abbas, S. I., Shah, M. H., & Othman, Y. H. (2021). Critical Review of Recruitment and Selection Methods: Understanding the Current Practices. *Annals of Contemporary Developments in Management & HR (ACDMHR)*, 3(3), Article 3. <https://doi.org/10.33166/ACDMHR.2021.03.005>
2. Abdollahpouri, H. (2019). Popularity Bias in Ranking and Recommendation. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 529–530. <https://doi.org/10.1145/3306618.3314309>
3. Abdollahpouri, H., & Mansoury, M. (2020). *Multi-sided Exposure Bias in Recommendation* (arXiv:2006.15772). arXiv. <https://doi.org/10.48550/arXiv.2006.15772>
4. *About Us*. (n.d.). Retrieved 12 September 2024, from <https://news.linkedin.com/about-us>
5. Ahmad, F. (2020). Personality traits as predictor of cognitive biases: Moderating role of risk-attitude. *Qualitative Research in Financial Markets*, 12(4), 465–484. <https://doi.org/10.1108/QRFM-10-2019-0123>
6. Åslund, O., & Skans, O. N. (2012). Do Anonymous Job Application Procedures Level the Playing Field? *ILR Review*, 65(1), 82–107. <https://doi.org/10.1177/001979391206500105>
7. Bartoš, V., Bauer, M., Chytilová, J., & Matějka, F. (2016). Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition. *American Economic Review*, 106(6), 1437–1475. <https://doi.org/10.1257/aer.2014.0571>
8. Behaghel, L., Crépon, B., & Le Barbanchon, T. (2015). Unintended Effects of Anonymous Résumés. *American Economic Journal: Applied Economics*, 7(3), 1–27. <https://doi.org/10.1257/app.2014.0185>
9. Bertrand, M., & Duflo, E. (2017). Field Experiments on Discrimination. In A. V. Banerjee & E. Duflo (Eds.), *Handbook of Economic Field Experiments* (Vol. 1, pp. 309–393). North-Holland. <https://doi.org/10.1016/bs.hefe.2016.08.004>
10. Bordalo, P., Coffman, K., Gennaioli, N., & Shleifer, A. (2016). Stereotypes\*. *The Quarterly Journal of Economics*, 131(4), 1753–1794. <https://doi.org/10.1093/qje/qjw029>
11. Bosak, J., & Sczesny, S. (2011). Gender Bias in Leader Selection? Evidence from a Hiring Simulation Study. *Sex Roles*, 9.
12. Brañas-Garza, P., Kujal, P., & Lenkei, B. (2019). Cognitive reflection test: Whom, how, when. *Journal of Behavioral and Experimental Economics*, 82, 101455. <https://doi.org/10.1016/j.socec.2019.101455>
13. Cauthen, N. R., Robinson, I. E., & Krauss, H. H. (1971). Stereotypes: A Review of the Literature 1926–1968. *The Journal of Social Psychology*. <https://www.tandfonline.com/doi/abs/10.1080/00224545.1971.9918526>
14. Ceschi, A., Costantini, A., Sartori, R., Weller, J., & Di Fabio, A. (2019). Dimensions of decision-making: An evidence-based classification of heuristics and biases. *Personality and Individual Differences*, 146, 188–200. <https://doi.org/10.1016/j.paid.2018.07.033>
15. Chen, J., Dong, H., Wang, X., Feng, F., Wang, M., & He, X. (2023). Bias and Debias in Recommender System: A Survey and Future Directions. *ACM Trans. Inf. Syst.*, 41(3), 67:1–67:39. <https://doi.org/10.1145/3564284>
16. Cho, J., Ahmed, S., Hilbert, M., Liu, B., & Luu, J. (2020). Do Search Algorithms Endanger Democracy? An Experimental Investigation of Algorithm Effects on Political Polarization. *Journal of Broadcasting & Electronic Media*. <https://www.tandfonline.com/doi/abs/10.1080/08838151.2020.1757365>
17. Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring Risk Literacy: The Berlin Numeracy Test. *Judgment and Decision Making*, 7(1), 25–47. <https://doi.org/10.1017/S1930297500001819>
18. Cole, M. S., Rubin, R. S., Feild, H. S., & Giles, W. F. (2007). Recruiters' Perceptions and Use of Applicant Résumé Information: Screening the Recent Graduate. *Applied Psychology*, 56(2), 319–343. <https://doi.org/10.1111/j.1464-0597.2007.00288.x>
19. Devine, P. G., Forscher, P. S., Austin, A. J., & Cox, W. T. L. (2012). Long-term reduction in implicit race bias: A prejudice habit-breaking intervention. *Journal of Experimental Social Psychology*, 48(6), 1267–1278. <https://doi.org/10.1016/j.jesp.2012.06.003>
20. Drezner, W., & Edigbe, E. (2024). *Accessible Low-Code No-Code Development*. <https://www.diva-portal.org/smash/get/diva2:1879862/FULLTEXT01.pdf>

21. D'Silva, C. (2020). A Study On Increase in E-Recruitment and Selection Process. *International Journal of Research in Engineering, Science and Management*, 3(8), Article 8.
22. Eliaz, K., & Schotter, A. (2010). Paying for confidence: An experimental study of the demand for non-instrumental information. *Games and Economic Behavior*, 70(2), 304–324. <https://doi.org/10.1016/j.geb.2010.01.006>
23. Ellemers, N. (2018). Gender Stereotypes. *Annual Review of Psychology*, 69(Volume 69, 2018), 275–298. <https://doi.org/10.1146/annurev-psych-122216-011719>
24. Eyting, M. (2022). Why do we Discriminate? The Role of Motivated Reasoning. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4210315>
25. Forscher, P. S., Mitamura, C., Dix, E. L., Cox, W. T. L., & Devine, P. G. (2017). Breaking the prejudice habit: Mechanisms, timecourse, and longevity. *Journal of Experimental Social Psychology*, 72, 133–146. <https://doi.org/10.1016/j.jesp.2017.04.009>
26. Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, 19(4), 25–42. <https://doi.org/10.1257/089533005775196732>
27. Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98(4), 506–528. <https://doi.org/10.1037/0033-295X.98.4.506>
28. Hosain, M. S., & Liu, P. (2020). The Role of Social Media on Talent Search and Acquisition: Evidence from Contemporary Literature. *Journal of Intercultural Management*, 12(1), 92–137. <https://doi.org/10.2478/joim-2020-0034>
29. Jowell, R., & Prescott-Clarke, P. (1970). Racial Discrimination and White-collar Workers in Britain. *Race*, 11(4), 397–417. <https://doi.org/10.1177/030639687001100401>
30. Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237.
31. Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature Human Behaviour*, 3(6), 587–595. <https://doi.org/10.1038/s41562-019-0589-3>
32. Koch, A. J., D'Mello, S. D., & Sackett, P. R. (2015). A meta-analysis of gender stereotypes and bias in experimental simulations of employment decision making. *Journal of Applied Psychology*, 100(1), 128–161. <https://doi.org/10.1037/a0036734>
33. Kordzadeh, N., & Ghasemaghahi, M. (2022). Algorithmic bias: Review, synthesis, and future research directions. *European Journal of Information Systems*, 31(3), 388–409. <https://doi.org/10.1080/0960085X.2021.1927212>
34. Kroll, E., Veit, S., & Ziegler, M. (2021). The Discriminatory Potential of Modern Recruitment Trends—A Mixed-Method Study From Germany. *Frontiers in Psychology*, 12. <https://www.frontiersin.org/articles/10.3389/fpsyg.2021.634376>
35. Kumar, V., Dudani, R., & K., L. (2023). The big five personality traits and psychological biases: An exploratory study. *Current Psychology*, 42(8), 6587–6597. <https://doi.org/10.1007/s12144-021-01999-8>
36. Lang, K., & Lehmann, J.-Y. K. (2012). Racial Discrimination in the Labor Market: Theory and Empirics. *Journal of Economic Literature*, 50(4), 959–1006. <https://doi.org/10.1257/jel.50.4.959>
37. *Linkedin recruiting trends (fr)*. (n.d.). Retrieved 14 September 2024, from [https://business.linkedin.com/content/dam/business/talent-solutions/regional/fr\\_FR/site/pdf/playbooks/linkedin-recruiting-trends-fr.pdf](https://business.linkedin.com/content/dam/business/talent-solutions/regional/fr_FR/site/pdf/playbooks/linkedin-recruiting-trends-fr.pdf)
38. Millecamp, M., Htun, N. N., Jin, Y., & Verbert, K. (2018). Controlling Spotify Recommendations: Effects of Personal Characteristics on Music Recommender User Interfaces. *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*, 101–109. <https://doi.org/10.1145/3209219.3209223>
39. Neumark, D. (2018). Experimental Research on Labor Market Discrimination. *Journal of Economic Literature*, 56(3), 799–866. <https://doi.org/10.1257/jel.20161309>
40. Pager, D., Bonikowski, B., & Western, B. (2009). Discrimination in a Low-Wage Labor Market: A Field Experiment. *American Sociological Review*, 74(5), 777–799. <https://doi.org/10.1177/000312240907400505>
41. Petit, P., Duguet, E., L'Hority, Y., du Parquet, L., & Sari, F. (2013). Discrimination à l'embauche: Les effets du genre et de l'origine se cumulent-ils systématiquement? *Economie et statistique*, 464(1), 141–153. <https://doi.org/10.3406/estat.2013.10234>

42. Roulin, N., & Levashina, J. (2019). LinkedIn as a new selection method: Psychometric properties and assessment approach. *Personnel Psychology*, 72(2), 187–211. <https://doi.org/10.1111/peps.12296>
43. Sczesny, S., Bosak, J., Neff, D., & Schyns, B. (2004). Gender Stereotypes and the Attribution of Leadership Traits: A Cross-Cultural Comparison. *Sex Roles*, 51(11–12), 631–645. <https://doi.org/10.1007/s11199-004-0715-0>
44. Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99. <https://doi.org/10.2307/1884852>
45. Singh, Y., Adil, Mohd., & Haque, S. M. I. (2023). Personality traits and behaviour biases: The moderating role of risk-tolerance. *Quality & Quantity*, 57(4), 3549–3573. <https://doi.org/10.1007/s11135-022-01516-4>
46. Statista. (2024). *Number of e-mail users worldwide 2027*. Statista. <https://www.statista.com/statistics/255080/number-of-e-mail-users-worldwide/>
47. Sue, S., & Kitano, H. H. L. (1973). Stereotypes as a Measure of Success. *Journal of Social Issues*, 29(2), 83–98. <https://doi.org/10.1111/j.1540-4560.1973.tb00074.x>