ROYAL HOLLOWAY UNIVERSITY OF LONDON

DOCTORAL THESIS

Explaining Human Oversampling Biases on Full Information Optimal Stopping Problems: a Behavioural, Computational and Neuroimaging Investigation

Author:
Sahira VAN DE WOUW

Supervisors:
Dr. Nicholas FURL
Prof. Ryan MCKAY

A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy

in the

Department of Psychology

January 28, 2022

Declaration of Authorship

I, Sahira VAN DE WOUW, declare that this thesis titled "Explaining Human Oversampling Biases on Full Information Optimal Stopping Problems: a Behavioural, Computational and Neuroimaging Investigation" and the work presented in it are my own. Any work done by myself jointly with others is declared here, along with my own contributions:

- **Chapter 1: Introduction**
  This is entirely my own work.
- **Chapter 2: Methodology**
  This is entirely my own work.
- **Chapter 3: Humans oversample in multiple decision-making domains**
  The study design was proposed by my supervisor, Dr Nicholas Furl, I supervised data collection, analysed the studies, and drafted the full manuscript.
- **Chapter 4: Explaining human sampling rates across different decision domains**
  I designed the studies, collected the data, analysed the studies, and drafted the full manuscript.
- **Chapter 5: Methodological remarks regarding optimal stopping tasks and the implications for sampling biases**
  I designed the studies, collected the data, analysed the studies, and drafted the full manuscript.
- **Chapter 6: Neural computations of prospective social decisions**
  I co-designed the study with my supervisor, Dr Nicholas Furl, collected the data, analysed the study, and drafted the full manuscript.
- **Chapter 7: Discussion**
  This is entirely my own work.

Signed:

Date: 28 January 2022

ROYAL HOLLOWAY UNIVERSITY OF LONDON

Abstract

Department of Psychology

Doctor of Philosophy

Explaining Human Oversampling Biases on Full Information Optimal Stopping Problems: a Behavioural, Computational and Neuroimaging Investigation

by Sahira VAN DE WOUW

An optimal stopping problem can be defined as a situation in which a decision-maker has to choose a time to take a given action. Within this thesis I look at a specific type of optimal stopping problem called the full information problem on which contrasting human behaviour has been reported. In full information problems, the decision-maker first learns the probability distribution that will generate the decision options, after which option values from this generating distribution are presented in sequence, and the decision-maker has to decide when to stop sampling and choose an option, under the condition that rejected options cannot be returned to later. The decision-makers' sampling rate is then compared to that of an optimal model to determine any sampling biases (undersampling or oversampling). My novel contribution to the literature is to show that human oversampling biases on these kinds of full information problems extend from the mate choice domain to other decision-making domains including image-based domains such as trustworthiness, foods and holiday destinations, as well as number-based domains such as smartphone prices. Furthermore, I describe how the moments of the generating distribution influence both the decision-makers' and the optimal model's sampling rate, and show that a correct specification of the generating distribution is crucial for correctly identifying sampling biases. Finally, I present neuroimaging evidence indicating that similar areas in the so-called decision network are activated when a decision-maker samples too few or too many options on a full information problem.

Acknowledgements

Firstly, I would like to thank my supervisory team, Dr Nicholas Furl, Prof Ryan McKay, and Prof Amina Memon. Nick has supported me from the very first Skype meeting we had back when I was still living in Amsterdam, and has continued to inspire me to become a better researcher. I especially would like to thank Nick for always having time to answer my (many) questions and to discuss any problems big or small. I could not have asked for a better lead supervisor. Ryan has a great eye for detail and has been invaluable at asking the critical questions and sharing insightful suggestions. I am very grateful for his input and invariably cheerful presence which made for some fun meetings. I would also like to thank Amina for encouraging me and always being willing to lend an ear whenever I needed it, you have been wonderful.

Secondly, I would like to thank the research assistants who have helped me with the data collection for this thesis: Santiago Torchiaro, Emanuel Zobev, Vrishni Kanhye, Jasneev Gill, and Gaia Giampietro. You have all been great. Also, a big thank you to my PhD cohort for the support and good times.

Thirdly, I would like to express my gratitude to my parents and my sister. Moving to the UK for this PhD project was a big decision and you have always been incredibly encouraging. Without your many greeting cards, care packages, and regular visits, I could not have accomplished this doctoral thesis.

And finally, I am eternally grateful that my research on the fiancé task has led me to my own fiancé. Callum Woods, thank you for always believing in me and being there for me, I could not have done this without you. I'm so lucky to have you by my side.

Contents

List of Figures

List of Tables

List of Abbreviations

ACC	**A**nterior **C**ingulate **C**ortex
AIC	**A**kaike **I**nformation **C**riterion
AP	**A**ttractive **P**rior
BF	**B**ayes **F**actor
BOLD	**B**lood **O**xygenation **L**evel-**D**ependent
BV	**B**iased **V**alues
C&A	**C**osta & **A**verbeck (2015)
CH	**Ch**apter
CI	**C**onfidence **I**nterval
CO	**C**ut**o**ff
dACC	**D**orsal **A**nterior **C**ingulate **C**ortex
DF	**D**egrees of **F**reedom
DLPFC	**D**orso**l**ateral **P**refrontal **C**ortex
DM	**D**ecision **M**aker
EEG	**E**lectro**e**ncephalo**g**ram
ERP	**E**vent-**R**elated **P**otential
fMRI	**f**unctional **M**agnetic **R**esonance **I**maging
FWE	**F**amily-**W**ise **E**rror
GB	**G**iga**b**yte
GBP	British Pound Sterling
HSO	Tukey's **H**onest **S**ignificant **D**ifference
ID	**Id**entity
IO	**I**deal **O**bserver
M	**M**ean
MDP	**M**arkov **D**ecision **P**rocess
MNI	**M**ontreal **N**eurological **I**nstitute
MRI	**M**agnetic **R**esonance **I**maging
OSF	**O**pen **S**cience **F**ramework
pgACC	**P**eri**g**enual **A**nterior **C**ingulate **C**ortex
RHUL	**R**oyal **H**olloway **U**niversity of **L**ondon
SD	**S**tandard **D**eviation
SR	**S**ample **R**eward
TE	**E**cho **T**ime
TR	**T**ime to **R**epeat
UK	**U**nited **K**ingdom
US	**U**nited **S**tates
VIF	**V**ariance **I**nflation **F**actor
VMPFC	**V**entro**m**edial **P**refrontal **C**ortex

Chapter 1

Introduction

Ever since their mysterious appearance in the mid-twentieth century, optimal stopping problems have been a popular field of study amongst scholars. To illustrate, a literature search in 2017 identified over 2,000 papers published on the classical secretary problem alone (Goldstein et al., 2017) (different variants are discussed later in this chapter). The topic has even made its way to the general public: an entire chapter was dedicated to optimal stopping problems in the popular science book "Algorithms to live by: The computer science of human decisions" (Christian & Griffiths, 2016) and multiple TED talks have been dedicated to optimal stopping scenarios (e.g., Fry, 2014; Griffiths, 2017). The popularity of optimal stopping problems is not entirely surprising. It has long been recognised that the type of sequential decision making involved in the class of optimal stopping problems this thesis focuses on is representative of many real-life decisions (Rapoport & Tversky, 1970). Accept a job offer or keep looking? Pick a parking spot or keep driving? Go on a date with your Tinder match or keep swiping right? For these kinds of decisions, we can ask ourselves: when is the optimal time to stop evaluating new information and commit to a decision?

The aim of this thesis is to understand sampling biases on full information optimal stopping problems. A sampling bias occurs when participants' sampling rate (i.e., how many options they sample before making a decision) differs from the optimal sampling rate. A distinction is made between oversampling biases (participants sample too many options) and undersampling biases (participants sample too few options). Note that sampling biases are relative to optimality. In other words, an increase in participants' sampling rate does not automatically mean that participants also oversample - the optimal sampling rate could have increased as well.

To understand sampling biases on full information problems, I take a holistic approach to decision-making by directly studying human behaviour as well as the neural network underlying decision-making. Additionally, I thoroughly examine the methodological approaches for studying optimal stopping problems, with the majority of chapters within this thesis investigating (amongst others) task features that could impact sampling biases. As such, a historical overview of the task designs used in previous research to study optimal stopping problems, as well as how they relate to the studies included in this thesis, falls within the scope of this Introduction. The Methodology chapter (Chapter 2) further builds upon this information and aims to clarify some of the methodological decisions made specifically for my studies.

Within this Introduction chapter I will first cover the theoretical background of optimal stopping problems, through which I will define the scope of this thesis (Section 1.1). Next, I will focus on human behaviour on optimal stopping tasks, and discuss sampling biases in Section 1.2. Different explanations for these observed sampling biases are explored in Section 1.3 and Section 1.4. In Section 1.5, I will discuss the neural correlates of decision-making, particularly on optimal stopping problems. At the end of each section, a paragraph is provided outlining how the information described in that particular section relates to the studies in this thesis ('Relevance'). The outline of the thesis and a summary of my research questions is provided in Section 1.6.

1.1 Types of optimal stopping problems

1.1.1 Classical secretary problem

The classical optimal stopping problem originated as a recreational mathematics problem called the secretary problem (for a summary of the historical background, see Ferguson, 1989; Freeman, 1983). The *classical secretary problem* is defined as follows: suppose your goal is to pick the highest ranking option from a set of options (e.g., the best candidate from a set of job applicants). The options are presented in series, one at a time, and the total number of options is known by the decision maker. The main constraint is that an option can only be chosen by the decision maker at the moment it is presented; it is not possible to go back to a past option (Chow et al., 1964; Rapoport & Tversky, 1970). The decision maker has no prior knowledge about what makes a good option or a bad option; they only know how the options compare to each other, that is, the relative ranks. This is why this classical example of the optimal stopping task is also referred to as the classic rank order version or the *no information problem* (Freeman, 1983; Gilbert & Mosteller, 1966; Petruccelli, 1980).

For mathematicians, solving the classical secretary problem meant proving that an optimal stopping rule existed (Chow & Robbins, 1963; DeGroot, 1968; Gilbert & Mosteller, 1966; Kahan et al., 1967; Shepp, 1969; Siegmund, 1967). Optimal stopping rules describe how individuals should make decisions on optimal stopping tasks in order to choose the best outcome. For the classical secretary problem, the optimal stopping rule can be described as

$$Z = \frac{N}{e} \tag{1.1}$$

with Z being the cutoff point, N the number of options, and e the base of the natural logarithm (equivalent to 2.71828...) (Kahan et al., 1967). The optimal strategy for the classical secretary problem is to keep sampling options, without choosing any, until Z is reached. During this time, some information about the distribution of option values can be accessed. The maximum value up until the cutoff point Z is remembered. After Z is reached, the next option that exceeds the remembered maximum is chosen as the optimal value. In practice, this converges to seeing at least $1/e$ = 37% of the total number of options before making a decision. As such, this optimal stopping rule is often referred to in later literature as the 37% rule (Ferguson, 1989). The success rate of the 37% rule, the number of times it results in choosing the best option, is also 37%.

1.1.2 Generalised secretary problem

The assumptions of the classical secretary problem described in Section 1.1.1 are rather restrictive, thus limiting its applicability to real-life decision making scenarios. Especially the assumption that only the highest ranking option yields a positive payoff is considered unrealistic. For many real-life decisions, alternatives other than the highest ranking option may still yield a positive payoff (Bearden et al., 2006). For example, hiring the second-best candidate from a set of job applicants still ensures that the position is filled and the job is done - a result that is generally considered as a positive outcome. The variant of the classical secretary problem that implements this adjustment is called the *generalised secretary problem*.

The generalised secretary problem is closely related to the classical secretary problem, but is less restrictive in the sense that participants can always expect a payoff, relative to the rank of the chosen option (Bearden & Connolly, 2007; Bearden & Murphy, 2007). Because the generalised secretary problem is considered to resemble real-life decision-making more closely, as decision makers do not only receive a payoff for choosing the best option but also for the second-best option, the third best option, etc., this variant is often preferred above the classical secretary problem. Bearden and Murphy (2007, p. 188) define the less restrictive condition as follows: "the [decision maker] earns a payoff of $\pi(a)$ for selecting an applicant with absolute rank a where $\pi(1) \geq ... \geq \pi(n)$". In this case, n is the number of options in the sequence. The optimal stopping rule remains the same as for the classical secretary problem, i.e., sample roughly 37% and then choose the next option which exceeds everything seen so far.

Another less restrictive version of the classical secretary problem, which closely resembles the generalised secretary problem, is the secretary problem with cardinal payoffs (Bearden, 2006). In the secretary problem with cardinal payoffs the decision maker receives a payoff equal to the option's underlying 'true' objective value (Bearden, 2006), rather than its relative rank as implemented in the generalised secretary problem (Bearden & Connolly, 2007; Bearden & Murphy, 2007). The optimal stopping rule for the secretary problem with cardinal payoffs is to skip the first $\sqrt{n} - 1$ options, after which the next option which exceeds everything seen so far should be chosen (Bearden, 2006).

Besides the generalised secretary problem and the secretary problem with cardinal payoffs, there are even more different variations of the classical secretary problem that have emerged over the years, all with different rules and assumptions (Freeman, 1983). Two of these variants - the partial information problem and the full information problem - will now be discussed in more detail.

1.1.3 Partial information problem

The term *partial information problem* captures a number of intermediate variants that fall between classical secretary (i.e., no information) problems and full information problems (Samuels, 2004). Before delving into what these variants are, I will first explain which characteristics of optimal stopping problems can be modified. Foremost, there is the distribution that generates the values of the options in a sequence. In other words, each option value in a sequence is drawn from a distribution of option values, also called the *generating distribution*. A common distribution on which optimal stopping tasks in previous experimental studies operate is a normal generating distribution, but there are some exceptions which I further discuss in Section 1.4.1. Key characteristics that can be modified in partial information problems are whether the participants know (or have learned) the mean value of the generating distribution, and whether the participants know (or have learned) the shape of the generating distribution (e.g., normally distributed). The following list includes examples of what could be called a partial information problem. Logically, a different optimal stopping rule is required for each variant, but to discuss them all would be beyond the scope of this thesis.

- Known generating distribution shape, unknown generating distribution mean (Sakaguchi, 1961).
- Known generating distribution shape, unknown number of options (Hill, 2009; Stewart, 1978).
- Known number of options, known generating distribution mean, unknown generating distribution shape (Hill, 2009).

1.1.4 Full information problem

One of the first to recognise the *full information problem* were Gilbert and Mosteller (1966). In full information problems, like (some) partial information problems but unlike secretary problems, the actual values of the options are presented rather than their relative ranks (Guan et al., 2014; Lee, 2006; Palley & Kremer, 2014; Shu, 2008). Also, there could be any number of mappings of rewards to ranks (which would be known to the participants) including, for example, a payoff relative to the rank of the chosen option, or equal to the option value itself. Furthermore, contrary to the classical no information secretary problem, the generating distribution of probabilities and values of all potential future options are known (Abdelaziz & Krichen, 2006;

Hill, 2009). Whereas in the classical secretary problem only the number of options was known, the full information problem also provides information about the continuous cumulative distribution. After each sample, the decision maker (who is familiar with the generating distribution and the number of options) is informed of the option value, and from this value can infer the probability of drawing any of the remaining values in future samples (Gilbert & Mosteller, 1966). This technique of inferring the probability of future values from previously drawn samples is called backward induction (see Sidebar 1; Hill, 2009), an idea stemming from game theory (Aumann, 1995).

Sidebar 1: Backward induction. In finite sequential search problems, like the full information problem, a participant who has arrived at the final option in a sequence can only take this option. Because of this property, it is possible to calculate the utility of every previous option by working backwards (i.e., using backward induction), starting with the final option using conditional probabilities (Costa & Averbeck, 2015). The decision threshold for taking the final option is ∞ because once a participant reaches the final option in a sequence, this option becomes their choice by default. Therefore, the utility of the final option is the participants' expectation of utility based on the known generating distribution (Baumann et al., 2020). This means that the decision threshold for the penultimate option is the expected utility of the final option. The utility for taking the penultimate option can then be calculated by multiplying the expected utility for this option with its probability plus the expected utility of the final option multiplied with its probability (Baumann et al., 2020). The utilities and decision thresholds for the remaining options can be calculated in the same way.

The optimal rule for decision making on full information problems is based on the idea of a variable threshold that is calculated for each position in the sequence (Gilbert & Mosteller, 1966). If an option value exceeds the threshold for that particular position in the sequence, the option is chosen. Because the decision maker is familiar with the generating distribution, they do not need the 37% rule to gather information about the distribution. As Gilbert and Mosteller (1966, p. 52) comment: "no buildup of experience is needed to set a standard, and a profitable choice can sometimes be made immediately". The trade-off between taking and declining an option on full information problems can be modeled using a Markov decision process (MDP). MDPs are an extension of the stochastic Markov chain model which describes a sequence of possible events. The difference is that MDPs not only model the utility (u) of a state (s) at sample (t), they also take into account actions and reward (Costa & Averbeck, 2015). As such, an MDP with only one available action and only one possible reward is in fact a Markov chain. MDPs can be applied to different variants of optimisation problems, including finite horizon, infinite horizon discounted, and infinite horizon average cost (Papadimitriou & Tsitsiklis, 1987). For full information problems, the utility of declining an option is effectively the decision threshold, as calculated using backward induction (see Sidebar 1).

1.1.5 Relevance

Section 1.1 describes four different types of optimal stopping problems: the classical secretary problem, the generalised secretary problem, the partial information problem, and the full information problem. The studies described within this thesis can all be classified as the latter. One of the reasons for this is that full information problems better resemble real-life optimal stopping scenarios than other versions of optimal stopping problems, because the assumptions are more realistic. For example, someone looking to buy a house will have some knowledge of the current market (i.e., familiarity with the underlying distribution from which options are sampled) and is likely to receive some payoff even if the house they end up buying was not the best they had seen (i.e., relative payoff). This application of full information problems to real-life decision-making situations has been demonstrated before, for instance for economic decision-making scenarios (e.g., "you are renting an apartment", "you are buying an ice cream sundae"; Cardinale et al., 2021; Costa & Averbeck, 2015), and even for a social mate choice scenario (Furl et al., 2019). As outlined in Section 1.1.4, full information problems are computationally particularly difficult to solve, which, combined with the application to real-life scenarios, makes it especially interesting and important to understand how participants solve such problems.

1.2 Decision biases

Having focused primarily on the theoretical background of optimal stopping problems in Section 1.1, I will now move on to human sampling behaviour on different optimal stopping tasks. As previously explained, there is an optimal solution to each task, which makes it possible to compare human performance to models of optimality. While the majority of early research on optimal stopping problems was mainly theoretical, Kahan et al. (1967) was one of the pioneers who investigated differences between human behaviour and mathematical optimality. In their study, 88 students were told to choose the highest number from a deck of 200 cards, which represented an offer for some stocks the participant was to sell. The results indicated that in this numerical version of the optimal stopping task, participants stopped sampling too early (*undersampling*), thus making suboptimal decisions.

At present, human undersampling biases are pervasive in the optimal stopping literature. This finding has been reported for different optimal stopping tasks, including the secretary problem (Bearden et al., 2006; Seale & Rapoport, 1997), numerical tasks (Guan et al., 2014; Kahan et al., 1967), the beads task (see Sidebar 2; Furl & Averbeck, 2011; Hauser et al., 2017; Hauser et al., 2018; Van der Leer et al., 2015), and tasks with different economic scenarios (Cardinale et al., 2021; Costa & Averbeck, 2015). Nevertheless, despite the overwhelming evidence for humans showing an undersampling bias, there are some conflicting findings. One recent paper by Furl et al. (2019), for example, showed consistent evidence for an *oversampling* bias in a mate choice scenario, which will be further discussed in Section 1.3. Several cognitive theories have been proposed in the literature that attempt to explain human sampling behaviour on optimal stopping tasks, a number of which are discussed below.

Sidebar 2: The beads task. An optimal stopping task that is closely related to the full information problem which I focus on, is the beads task (Furl & Averbeck, 2011; Van der Leer et al., 2015). The beads task is also known in the literature as the 'information sampling task' (Hauser et al., 2017; Hauser et al., 2018) or the 'probabilistic inference task' (Huq et al., 1988), but I will use the term beads task and describe it as such. In the beads task, participants are presented with sequences of coloured beads drawn from a hidden urn. The goal is to correctly guess the majority bead colour of the urn. After each bead, participants have two options: either to draw another bead, or to attempt to guess the urn's contents. The beads task relates to the full information problem in that participants are presented with a similar decision-making scenario, namely, either to continue sampling or to stop sampling and make a decision. Furthermore, participants are often informed of the proportion of bead colours in the urn prior to the task (e.g., 60/40; Furl & Averbeck, 2011), and there is a finite number of beads in each sequence (e.g., 25; Hauser et al., 2017). Research on the beads task generally reports that participants sample too few options compared to models of optimality (Furl & Averbeck, 2011; Hauser et al., 2017; Hauser et al., 2018; Van der Leer et al., 2015). Mathematically, the optimal model for the beads task is very similar to the optimal model that can be used to solve full information problems (Cardinale et al., 2021; Costa & Averbeck, 2015), in that the model uses Bayes' rule (see Section 2.1.1) and backwards induction techniques (see Sidebar 1).

1.2.1 Overestimation of positive payoffs

Bearden et al. (2006) studied the generalised secretary problem and described two experiments in their paper. After finding in their first experiment that participants terminated their search too early, Bearden et al. (2006) conducted a second experiment to answer the question of why participants undersample. Thirty participants completed a generalised secretary task where only relative ranks were shown and payoffs decreased linearly as the quality of the selected applicant decreased. Recall that this means that participants get paid for selecting any option, not just for selecting the best (classical secretary problem). After the secretary task, participants completed a probability estimation task on which they were asked to estimate the probability that an applicant's absolute rank was the same as the applicant's relative rank (for an example, see Table 1.1). Participants provided responses using a slider from 0 to 100, and completed a total of 72 probability estimates. The results of this second experiment revealed that participants tend to overestimate the true probabilities. In other words, they believe the absolute rank, and thus their payoff, is greater than it actually is based on the relative ranks they saw. Therefore, Bearden et al. (2006) hypothesise that the undersampling bias on rank-based optimal stopping tasks is a result of participants overestimating the quality of options. One limitation of this study is that the results cannot be applied to variants of optimal stopping tasks that display the actual value of an option or the absolute rank, rather than an option's relative rank (i.e., full information problems). The reason for this is that in these kinds of full information problems, participants can infer their payoff directly from the observed option value rather than having to estimate their payoff from the options' relative ranks, leaving little room for the overestimation of positive payoffs.

TABLE 1.1: A visualisation of absolute ranks and relative ranks in a sequence with five options.

Position in the sequence	1	2	3	4	5
Absolute rank	2	5	3	1	4
Relative rank	1	2	2	1	4

The principle on which Bearden et al. (2006)'s hypothesis is based, however, finds support in other sequential search tasks. Mantonakis et al. (2009), for example, described the primacy effect, or first-impression effect, which was found to influence performance on a sequential search task. The primacy effect describes how options early in a sequence leave a bigger impression than options later on in a sequence, thus giving early options an advantage. It should be noted that Mantonakis and colleagues used a sequential search task where a single choice of the best option was made at the end of the sequence. As such, this was not strictly an optimal stopping problem. Nevertheless, the primacy effect is in line with the hypothesis proposed by Bearden et al. (2006), and could explain how early options on an optimal stopping task seem over-proportionately better compared to later options causing participants to terminate their search too early.

1.2.2 Learning effects

Another factor that could potentially influence human sampling rate on optimal stopping tasks is learning, in the sense of participants using knowledge of sampled values from preceding sequences to make decisions on subsequent sequences. Goldstein et al. (2020), for example, reported that if the underlying distribution is unknown, participants can correct an initial undersampling or oversampling bias with repeated play. On this classical secretary problem, fast and steep learning effects were observed. However, contrasting results have been found for full information problems. Campbell and Lee

(2006), for example, investigated whether people improve over repeated trials of a full information problem, but concluded that there was no evidence of learning. Similarly, Lee (2006) observed a lack of learning on another full information problem, and concluded that models of human behaviour do not need to incorporate learning.

To verify that learning indeed does not have an effect on full information problems, I examined learning effects in two of my studies (Study 1 and Study 2 as described in Chapter 4). Participants' sampling rate is plotted as a function of sequence number (Figures 1.1 and 1.2). As expected, sequence number appears to have a minimal effect on participants' sampling rate. From this I conclude that learning effects are unlikely to have an influence on the findings presented in this thesis.

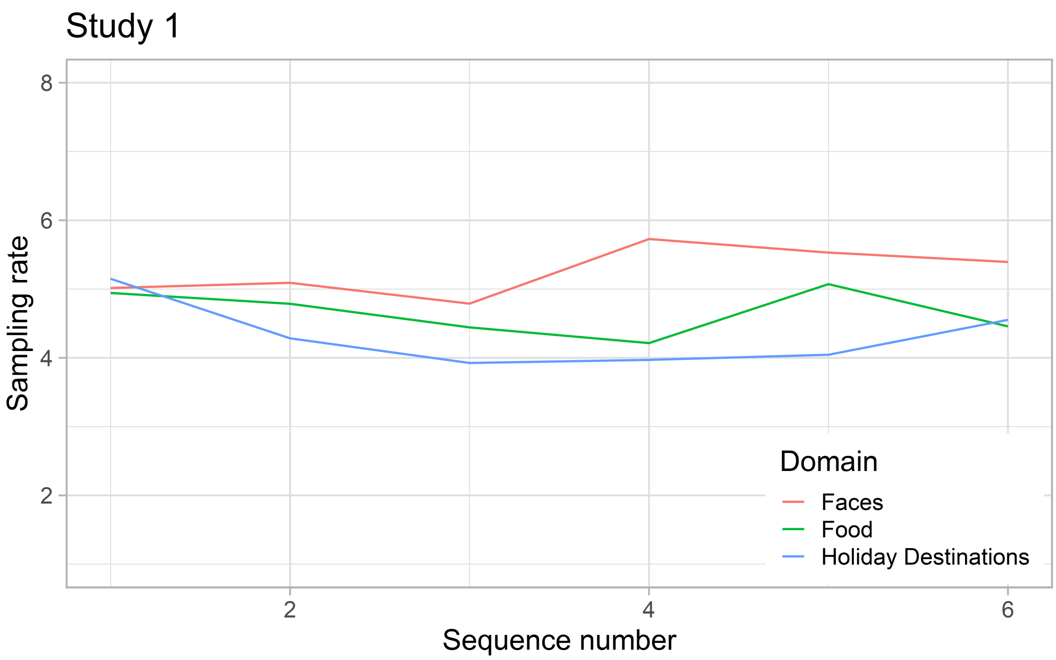

FIGURE 1.1: Participants' sampling rate plotted as a function of sequence number. Study 1 as described in Chapter 4 includes three decision-making domains, which are depicted here (Faces, Food and Holiday Destinations). Minimal learning effects can be observed across all three domains.

1.2.3 Relevance

Neither overestimation of positive payoffs nor learning effects are expected to affect the studies included in this thesis. Overestimation of positive payoffs is limited to studies that display options' relative ranks, and none of the studies included here can be classified as such because they are all full

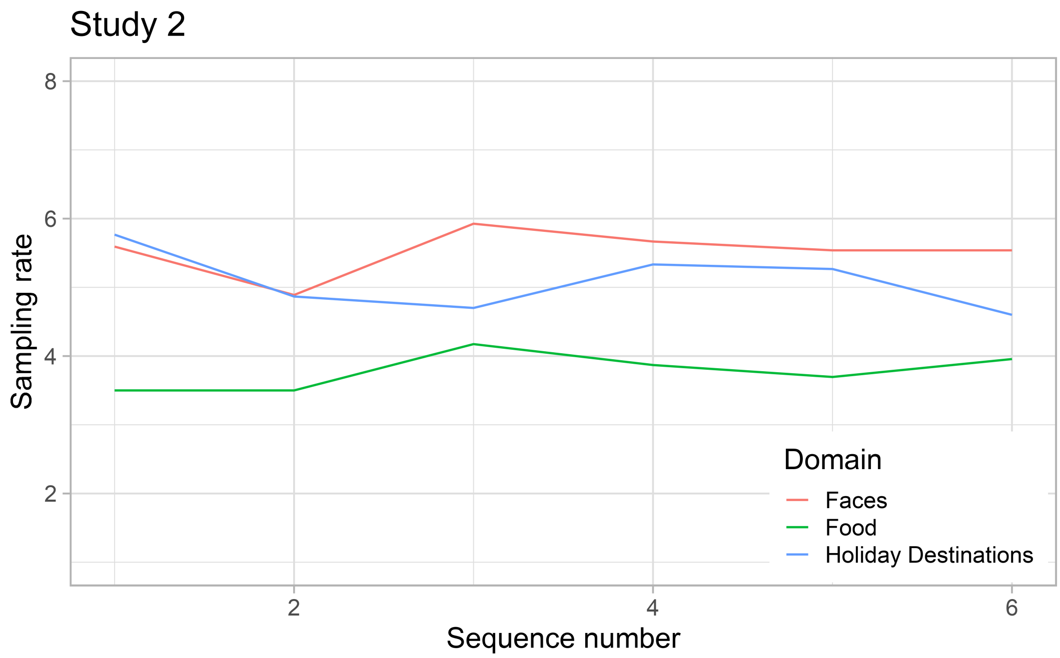

FIGURE 1.2: Participants' sampling rate plotted as a function of sequence number. Study 2 as described in Chapter 4 includes three decision-making domains, which are depicted here (Faces, Food and Holiday Destinations). Minimal learning effects can be observed across all three domains.

information problems. Learning effects have not been observed in full information problems either, a finding which I have confirmed using my data of Chapter 4.

1.3 Decision-making domain

In Section 1.2, I briefly mention the oversampling bias reported by Furl et al. (2019). Furl et al. (2019) studied mate choice decisions on a full information problem dubbed the facial attractiveness task. In phase 1 of the facial attractiveness task, participants were instructed to rate images of faces (of their preferred partner's sex) on their attractiveness. In phase 2, participants encountered multiple sequences in which they attempted to maximise the attractiveness of their date. Furl et al. (2019) conducted three separate studies, varying the number of sequences (from 5 to 28) and the length of the sequences (8 or 12), but found similar results: participants sampled more and ended up with lower-ranked faces compared to a model of optimality. In other words, humans show an oversampling bias on the facial attractiveness task. The authors propose that the oversampling bias might be specific to the mate choice domain (Furl et al., 2019).

Furl et al. (2019) hypothesised that the design of the task was sufficient to instigate mate choice predispositions, i.e., the tendency to set high thresholds and continue searching for a high quality partner. Upon closer inspection, participants in Furl et al. (2019) were observed to keep their decision thresholds too high throughout the sequences. This could have led to the observed oversampling bias on this particular full information problem. These kind of mate choice predispositions are not limited to humans, but have been observed in the animal literature as well, e.g., in crickets (Ivy & Sakaluk, 2007), fiddler crabs (Backwell & Passmore, 1996), and sticklebacks (Milinski & Bakker, 1992).

Traditional mathematical treatments of optimal stopping problems have anecdotally described the problem in mate choice terms such as 'fiance' or 'dowry' problem (Gilbert & Mosteller, 1966). Yet the problem was not framed as mate choice in an empirical study in humans until Furl et al. (2019), who have created a contemporary paradigm that captures the essence of an online dating scenario (e.g., Tinder). This differs from other mate choice paradigms that have attempted to incorporate 'real-life' elements such as rejection (Miller & Todd, 1998), interaction (Eriksson & Strimling, 2009) and self-perceived attractiveness (Beckage et al., 2009) in their tasks. Regardless of whether the facial attractiveness paradigm embodies every element that can occur in real-life dating, it is still possible that the mate choice domain was instrumental in changing participants' search strategies from undersampling to oversampling.

1.3.1 Relevance

Before this thesis, there was no conclusive evidence that could explain the oversampling bias reported on the facial attractiveness task (Furl et al., 2019). The leading hypothesis was that the decision-making domain has influenced participants' sampling behaviour. Whether oversampling biases are indeed specific to the mate choice domain is one of the questions I aim to address in the first two experimental chapters of this thesis (Chapters 3 and 4), thus marking the start of my exploration into what causes participants' sampling biases on full information optimal stopping tasks.

As announced in Section 1.1, all studies described in this thesis can be classified as full information problems. The reason for this is that one of the main aims of this thesis is to understand why Furl et al. (2019) found oversampling biases on their full information facial attractiveness task. Real-life mate choice decisions could include many elements of full information problems, and because of this realistic aspect, full information problems are especially worthy of study. To be able to compare my findings to those of previous research (Furl et al., 2019), I too employ full information tasks, and define my scope as such to only include full information problems.

1.4 Task features

Thus far I have discussed cognitive theories that are thought to affect sampling biases, and I have discussed literature that suggests that the decision-making domain can influence participants' sampling behaviour. The final piece of information to understand sampling biases on optimal stopping problems might be the most extensively researched one: the effect of task features. As Goldstein et al. (2020) observed, minor differences in instructions to participants can have major effects on the results. Within this section, I will discuss a selection of task features that have been found to influence participants' sampling rate.

1.4.1 Varying generating distributions

Guan et al. (2014) investigated the effect of varying the nature of the generating distribution of option values on participants' choice thresholds, in a partial information problem where participants knew the number of options, but not the shape or mean of the generating distribution. The only way the generating distribution could be learned, was through sampling. Fifty-six participants were evenly divided between two conditions. Condition one was a scarce environment, meaning that the generating distribution consisted mainly of low values. Condition two was a plentiful environment, meaning that the generating distribution consisted mainly of high values. The optimal threshold in the plentiful environment is higher than the optimal threshold in the scarce environment. Guan et al. (2014) define two variables of interest: correspondence and coherence. Correspondence refers to how often participants choose the best value in a sequence. Coherence refers to how often participants make decisions that are the same as the optimal model.

Additionally, participants' thresholds were computed using Bayesian methods to fit threshold models to all of the individual participants. Guan et al. (2014) found that participants updated their thresholds in accordance with the distributional properties of their environment, and participants' performance (i.e., correspondence and coherence) was relatively good.

The aim of Guan et al. (2014) to investigate the effect of varying generating distributions on sampling behaviour is not a novel one. Kahan et al. (1967) already hypothesised that distributions differing in variance and skewness might lead to differences in sampling behaviour. Three distributions with the same mean but different standard deviations were compared: a positively skewed triangular distribution (A), a negatively skewed triangular distribution (B), and a rectangular distribution (C). Participants could only learn these distributions through sampling, like in Guan et al. (2014), therefore the paradigm of Kahan et al. (1967) could be classified as a partial information problem too. Kahan et al. (1967), however, found insufficient evidence to support their hypothesis, as participants' mean number of samples did not differ significantly between the three distributions A, B, and C. Both in Guan et al. (2014) and Kahan et al. (1967), the tendency to stop sampling too early was maintained in all environments.

A more recent paper by Baumann et al. (2020) examined the effect of varying generating distributions of option values on a full information problem. The second experiment in Baumann et al. (2020) investigates human behaviour in three different task environments: a scarce environment (left-skewed distribution), a normal environment (normal distribution), and a plentiful environment (right-skewed distribution). The experiment included a learning phase, where participants learned from which distribution options were sampled (through descriptions using statistical terminology and graphs of the probability densities of statistical distributions), and a testing phase, where participants had to choose the lowest-priced ticket from a sequence of 10 ticket prices. The results showed that the mean number of samples for participants decreased from the left-skewed environment to the right-skewed environment, thus indicating that participants' sampling rates are affected by the skewness of the generating distribution (Baumann et al., 2020). This is in contrast with the findings of Kahan et al. (1967) who did not report a difference in sampling rate. However, Kahan et al. (1967) employed a partial information problem and Baumann et al. (2020)'s paradigm was a full information problem. Also, I note that the generating distributions used by

Baumann et al. (2020) to populate their scarce and plentiful environments did not just differ in skewness, but in mean as well. This highlights the need for further research into how the moments of the generating distribution (i.e., mean, skewness, kurtosis and variance) affect human sampling behaviour: a question I will address in Chapter 4.

1.4.2 Time cost

In Sections 1.2 and 1.4.1 I have discussed the study by Kahan et al. (1967), who reported undersampling on a task where participants had to choose the highest number from a deck of cards. Interestingly, the same participants tended to sample too much when tested in groups, rejecting options that, based on the optimal model, should have been accepted. The hypothesis proposed by Kahan and colleagues prior to the experiment was that the time participants have to invest in the experiment is a variable factor that determines how costly it is for participants to keep sampling options. Henceforth, I will refer to this particular cost-to-sample as *time cost*. Logically, if the time cost is high, participants will sample less, and if the time cost is low, participants will feel free to sample more. When participants were tested individually, they were free to leave once they finished the experiment. In other words, the time cost is relatively high when tested individually, as participants have to invest their free time to sample more. In the group condition, participants had to stay until the slowest member of the group had finished. Therefore, the time cost is lower when tested as a group because the participants did not lose any time sampling more options. The undersampling and oversampling biases reported by Kahan et al. (1967) for participants in the individual and group conditions, respectively, seems to support the hypothesis that time cost can limit how much participants sample.

Three decades after the study by Kahan et al. (1967) was published, Seale and Rapoport (1997) replicated Kahan et al. (1967)'s undersampling findings in a classical secretary problem. Recall that the optimal stopping rule for classical secretary problems is the 37% rule, which describes the optimal cutoff point after which the next option that exceeds the maximum value seen so far is chosen. In line with Kahan et al. (1967), Seale and Rapoport (1997) hypothesised that the undersampling bias observed in participants could be attributed to the existence of an endogenous time cost. Using computer simulations, Seale and Rapoport (1997) were able to show that introducing a search cost as low as 0.002 effectively reduces the optimal cutoff point. This means that if participants factored time cost in their sampling behaviour, it would result in earlier termination of searches.

A study that looked at endogenous search costs in a full information problem is Costa and Averbeck (2015). In Costa and Averbeck (2015), cost-to-sample was a free parameter in their optimal model which could be adjusted to maximise the fit of the optimal model to participants' data. Costa and Averbeck (2015) reported that a positive cost-to-sample value significantly improved the optimal model's predictions of human behaviour, as it effectively decreased the value of future samples for the model, leading to earlier termination of searches. In the absence of any extrinsic costs associated with the task, this result suggests that participants made decisions with an endogenous search cost, i.e., time cost, in mind.

The research discussed in this section suggests that time cost could explain human undersampling biases on classical secretary problems (Seale & Rapoport, 1997), partial information problems (Kahan et al., 1967), and full information problems (Costa & Averbeck, 2015). This evidence raises the question, if a positive cost-to-sample can explain undersampling, could a negative cost-to-sample explain oversampling? This question has been investigated by Furl et al. (2019), who compared participants' sampling rate to that of a Bayesian computation model (see Chapter 2) with a negative cost-to-sample parameter, that is, the model assigned a reward value to sampling more options. This is effectively the opposite of implementing a time cost (i.e., a positive cost-to-sample) to explain undersampling biases (Costa & Averbeck, 2015). Yet, Furl et al. (2019) found that the choice thresholds produced by the 'sample reward model' did not match participants' thresholds. This suggests that oversampling biases cannot be explained by an endogenous reward for sampling.

1.4.3 Monetary cost

The facial attractiveness task (as discussed in Section 1.3) may have been unique, but Furl et al. (2019) were not the only researchers to report that participants show an oversampling bias on an optimal stopping task. On a generalised secretary problem (see Section 1.1) where recall of previously rejected options was allowed (see Sidebar 3), Zwick et al. (2003) reported that participants oversampled when there was a fixed monetary cost-to-sample. To avoid confusion with the previously defined cost-to-sample (time cost, see

Section 1.4.2), I will henceforth refer to a monetary cost-to-sample as *monetary cost*. The task employed by Zwick et al. (2003) was framed as an economic scenario where participants attempted to select the best apartment to rent. Compared to a condition without a monetary cost, participants sampled less in the condition with a monetary cost. In other words, participants' sampling rate decreased when a monetary cost was added. However, the optimal model decreased its sampling more than participants in the monetary cost condition, resulting in participants showing an oversampling bias. In this thesis, I consider only the influence of time cost on sampling biases (see Section 1.4.2), and not monetary cost, as none of the studies presented here include such an extrinsic cost-to-sample.

Sidebar 3: Recall. The classical secretary problem does not allow for participants to go back and choose a previously rejected option (i.e., recall). However, there are certainly situations in everyday life where recall is an option. A typical example in the field of optimal stopping problems of a case where recall was possible is the story of how famous astronomer Johannes Kepler (1571-1630) set out to find his second wife (Ferguson, 1989). After interviewing eleven candidates and rejecting all, he returned to candidate number five, won her over, and married her. The possibility of recall was recognised early on, amongst others by DeGroot (1968). DeGroot reasoned that sampling with recall only extends an additional advantage on partial information problems, and not on full information problems. If a certain option on a full information problem was not considered good enough the first time around, i.e., it did not pass the fixed threshold, it is never going to be good enough and should therefore never be recalled. The fact that the possibility of recall should not have an effect on sampling behaviour on the full information problem does not mean that it has no effect. Indeed, Kogut (1990) reported that if a monetary cost-to-sample is added to the full information problem, recall increases to 42.6% of the time, compared to 11% of the time when there is no monetary cost-to-sample, despite the irrelevance of the recall option in both variants (monetary cost/no monetary cost). On a partial information problem, however, where sampling is needed to learn more about the true state of the environment, options that seemed insufficient at an early stage could become acceptable at a later stage if the true state of the environment is not as favourable as the decision maker thought it was (DeGroot, 1968).

1.4.4 Reward structure

There are conflicting findings regarding the effect of reward structure on sampling behaviour. With the term 'reward structure' I refer to the way participants are compensated for their participation in a study, including any additional (extrinsic) rewards that may be obtained by participants along the way (e.g., bonus payments). A study by Hsiao and Kemp (2020), described the effect of different reward structures on search behaviour in two secretary problems. The results indicated that when participants were rewarded only for obtaining the best option in the sequence (one of the assumptions of the classical secretary problem), they searched longer in both paradigms compared to the 'commission base' and 'flat fee' reward structures. Furthermore, participants in the 'best only' group more frequently obtained the best option, indicating that their search behaviour was more optimal.

A study by Campbell and Lee (2006), however, found that on a full information problem, financial reward by itself did not impact performance. The two possible reward structures in Campbell and Lee (2006) were 'no extrinsic reward' and a 'quota-piece rate scheme' where participants were rewarded for high performance. In the latter, participants were rewarded an additional \$5 after every 12 correct responses, on top of a \$5 flat fee, with a ceiling of \$30. However, Campbell and Lee (2006) did not look at the effects on sampling rates, but rather on learning (improvement over trials) and overall performance (proportion of correct decisions). Neither was found to be affected by financial reward.

The majority of the studies included in this thesis compensate participants for their time with a flat fee, and do not incorporate any additional reward structures that might affect performance. Participants who are paid a flat fee are nevertheless instructed to try to obtain certain options, e.g., as attractive an option as possible. In obtaining these, participants presumably experience some kind of intrinsic reward, in line with performing as they are instructed, that motivates the strategies they use. Chapter 5 makes a humble contribution to the literature on reward structures by comparing a flat fee reward structure to a reward structure where participants receive bonus payments on top of a flat fee for obtaining the best, second best, or third best option in the sequence (as implemented in Costa & Averbeck, 2015).

1.4.5 Previous values

Lee (2006) investigated whether a participant's choice to stop sampling on a full information problem is sensitive to previous options that were seen. However, the results showed that there was no such relationship between rejected options and the chosen option. Neither the immediately preceding option nor any other preceding options were predictive of whether or not an option would be chosen. Therefore, Lee (2006) concludes that models of human behaviour on optimal stopping tasks do not need to be sensitive to previous values in a sequence to explain performance. This finding was later confirmed in Guan et al. (2014) and Guan and Lee (2018), who found that participants' choice thresholds were independent of the preceding value in a particular sequence, meaning that participants did not make decisions based on the context provided by earlier options. Similar results have been observed on a classical secretary problem. Seale and Rapoport (2000) investigated two models where the agent keeps count of either high-ranked or low-ranked options, and only makes a decision after a certain number of high-ranked or low-ranked options has been seen. However, neither model was found to be a good description of human decision-making with the majority of participants (84%) adhering to a threshold model instead (the cutoff model; see Chapter 2, Section 2.1.3 for a more detailed description). Following the evidence presented here that previous values in a sequence do not significantly affect whether the next option is chosen, I will not further touch upon this topic in this thesis.

1.4.6 Relevance

To summarise, in this thesis I will further discuss the effect of varying generating distributions. Specifically, Chapter 4 includes a thorough investigation of the moments of the generating distribution and their influence on participants' sampling rate. Furthermore, Chapter 5 contributes to the research on reward structure by examining whether awarding bonus payments for the three highest ranking options could affect participants' sampling rate. Also addressed in Chapter 5 is time cost. Specifically, I compare an optimal stopping task with fixed timings to a self-paced optimal stopping task to see if this has an effect on participants' sampling bias. Time cost is not further addressed in any of my other experimental chapters as they aim to explain human oversampling biases on full information problems. Previously research has found that time cost was only able to explain undersampling biases (Costa & Averbeck, 2015; Kahan et al., 1967; Seale & Rapoport, 1997), which logically makes sense. If participants sample with an endogenous cost of search in mind, sampling more options than optimal would reduce the participants' payoff per unit time (Seale & Rapoport, 2000). As mentioned before in Section 1.4.3, none of the studies presented in this thesis include an extrinsic monetary cost-to-sample which is why I will not further discuss this task feature, nor will I focus on previous values as these cannot explain participants' performance (Lee, 2006).

1.5 Neural correlates of decision-making on optimal stopping tasks

Fundamental to understanding sampling biases on optimal stopping problems is to delineate the underlying neural correlates of decision-making on these kind of tasks. Optimal stopping tasks require participants to make prospective decisions where they weigh current rewards against prospective probability of future reward. In other words, participants must either take the current option or decline it with prospective expectations of a better future option. Previous research has linked several areas to prospective decision making, including the dorsal anterior cingulate cortex (dACC), dorsolateral prefrontal cortex (DLPFC), posterior cingulate cortex and perigenual anterior cingulate cortex (pgACC) (Kolling et al., 2016a; Kolling et al., 2018; Kolling et al., 2016b). However, these studies do not incorporate evidence seeking (i.e., sampling of options to gain information) in their task design. For optimal stopping tasks, evidence seeking can be considered its principal feature (Kolling et al., 2018).

One of the key papers that discusses the neural correlates of prospective decision-making on optimal stopping problems is Costa and Averbeck (2015). In their study, participants made decisions on an economic optimal stopping task (e.g., 'find a car with the lowest mileage') and were rewarded for picking one of top three highest ranking options. Option lists were drawn from 14 different categories (e.g., buying a subway ticket, a television, or renting an apartment), and sequences comprised either 8 or 12 options. Participants' sampling behaviour was compared to that of a Bayesian ideal observer model, and results showed that participants sampled too little compared to the optimal model (Costa & Averbeck, 2015). More interestingly, when examining the contrast between choices to take an option versus choices to decline an option, a robustly activated set of areas was revealed:

- Anterior insula bilaterally
- Dorsal anterior cingulate cortex
- Ventral striatum
- Parietal-frontal areas - specifically left lateralised dorsal parietal cortex and left lateralised prefrontal cortex

The reason Costa and Averbeck (2015) looked specifically at the contrast between take versus decline choices was because this particular contrast resembles the crossing of the choice threshold (see Section 1.1.4). A contrast between take versus decline choices (take > decline) should be interpreted as the blood oxygenation level-dependent (BOLD) response that is unique for choices to take the current option. The take versus decline contrast is achieved by subtracting the BOLD response for decline choices from the BOLD response for take choices (take - decline). Costa and Averbeck (2015) also examined the opposite contrast, decline versus take choices (decline > take), but found no significant clusters of voxels for this contrast.

The areas listed above are largely in alignment with other studies investigating the neural correlates of take versus decline decisions, as implemented in optimal stopping problems. Furl and Averbeck (2011), for example, looked at similar prospective decisions in a related paradigm called the beads task (see Sidebar 2). In both Costa and Averbeck (2015) and Furl and Averbeck (2011), the contrast compares the choice to take an option (and stop sampling) versus the choice to decline an option (and continue sampling). Like Costa and Averbeck (2015), Furl and Averbeck (2011) reported activation in the anterior insula, anterior cingulate, ventral striatum, and parietal cortex. Background information on the function(s) of the regions related to prospective decision-making is provided below.

1.5.1 Anterior insula

The anterior insula is a bilateral cortical region that has many different functions, including in decision-making (Furl & Averbeck, 2011; Loued-Khenissi et al., 2020). For example, previous research has found that the insula is involved in processes like conscious error perception (Klein et al., 2013), uncertainty (Huettel et al., 2005; Loued-Khenissi et al., 2020), and evaluating stimulus significance (Thielscher & Pessoa, 2007). In the context of optimal stopping problems (the beads task), activation in the insula was found to be modulated by potential losses and expected gain (Furl & Averbeck, 2011). This means that the size of the BOLD response in the insula covaried with the action value Q (i.e., the expected gain), which corresponds to the action (take or decline) chosen by the participant for each option. As such, the anterior insula is thought to relate to the value of the current option, thus contributing to the termination of evidence seeking, along with other areas as discussed below (Furl & Averbeck, 2011).

1.5.2 Anterior cingulate cortex

Research in macaque monkeys has demonstrated that the ACC plays a vital role in learning and maintaining the ongoing value of actions (Kennerley et al., 2006). One of the ways in which the ACC is thought to compute action values is through an integration of decision variables including action and reward history, risk, and expected payoff (Kennerley et al., 2006). This is relevant to optimal stopping problems, as calculating the value of either taking or declining an option in a sequence is central to optimal decision making. There is a robust body of research in humans supporting these findings. For example, results on a patch foraging paradigm - a related decision-making task - reported that humans retain the values of past choices as well as the prediction of future choice values in the dACC (Wittmann et al., 2016). Moreover, not only does the dACC hold past and future choice values, it encodes the value of explorations, i.e., search values, as well (Kolling et al., 2016b). All in all, the ACC appears to play a crucial role in evaluating potential future actions (Schuck et al., 2015; Wittmann et al., 2016), which is a key component of optimal stopping tasks. In addition to its role in computing action values, the ACC has also been found to have a function in response conflict (Knutson et al., 2007; Ridderinkhof et al., 2004) and reward (Haber & Knutson, 2010). Kuhnen and Knutson (2005), for example, found activation in the ACC related to gain outcomes on a financial decision-making task, particularly under conditions of increased response conflict.

1.5.3 Ventral striatum

The ventral striatum is a key structure that responds to reward (Haber & Knutson, 2010). Specifically, activation in the ventral striatum is associated with the anticipation of rewards, i.e., monetary gain (Haber & Knutson, 2010; Jauhar et al., 2021; Knutson et al., 2001a; Knutson et al., 2001b). This role of the ventral striatum in reward on optimal stopping tasks has been confirmed by Costa and Averbeck (2015), who reported that neural responses in the ventral striatum were parametrically modulated by the size of the reward outcome. An alternative hypothesis might be that activation in the ventral striatum could be related to the anticipation of the reward feedback screen in the take versus decline contrast. In Costa and Averbeck (2015), when a decision maker has decided to take an option, they are directed to a reward feedback screen showing them whether they had won $5, $3, or $1 (or nothing). It is possible that the ventral striatum played a role in the anticipation of this reward feedback screen after an option had been chosen.

1.5.4 Parietal cortex

For take versus decline choices, both Costa and Averbeck (2015) and Furl and Averbeck (2011) report peak activation in the inferior parietal lobule. However, the parietal cortex comprises many other sub-regions. Aside from the inferior parietal lobule, I here also discuss the posterior parietal cortex as this sub-region seems involved in processes relevant to prospective decision-making.

The parietal cortex plays an important role in selective attention (Kastner & Ungerleider, 2000) and decision making. Studies in macaque monkeys, for example, have shown that planning movements to make a free choice between targets activated the parietal cortex as well as the frontal cortex (Pesaran et al., 2008; Platt & Glimcher, 1999). Furthermore, Platt and Glimcher (1999) found that the gain expected from each possible action (i.e., the action values) correlated with activation of posterior parietal neurons, which is in line with research suggesting that the parietal cortex is involved in the integration of evidence as it is collected (Kiani & Shadlen, 2009). In fact, there is considerable support for the parietal cortex, and particularly the lateral intraparietal cortex, to be involved in evidence accumulation in monkeys (Beck et al., 2008; Shadlen & Newsome, 1996). Despite the role of the posterior parietal cortex in evidence accumulation, this area has not been previously linked to decision-making on optimal stopping tasks.

Instead, Furl and Averbeck (2011) found that on the beads task, activation in the inferior parietal lobule was associated with decision thresholds and sampling rate. Specifically, Furl and Averbeck (2011) reported that for take > decline decisions, the size of the BOLD response in the parietal cortex showed a significant linear relationship with individual differences in the participants' mean number of samples, with parietal responses being greater for participants' who sampled more options. Furthermore, an electroencephalogram (EEG) study by Kopp et al. (2016) on the beads task showed that manipulations of prior probabilities and likelihoods (i.e., Bayesian inference) were associated with amplitude variations in event-related potentials (ERPs) in the frontal cortex and parietal cortex, respectively. Not only does this study demonstrate the capability of the human brain for Bayesian inference, it highlights the role of the parietal cortex as well as the frontal cortex in decision-making on the beads task (which, as we know, is a type of optimal stopping problem).

1.5.5 Prefrontal cortex

Like the parietal cortex, the prefrontal cortex comprises a large region consisting of many sub-regions. Costa and Averbeck (2015) report activity in the DLPFC as well as the ventromedial prefrontal cortex (VMPFC), and these are therefore the two sub-regions I will focus on here.

The previous section already highlighted that the prefrontal cortex is co-activated with the parietal cortex in a decision-making task in monkeys (Pesaran et al., 2008), as well as in the computation of prior probabilities and likelihoods on the beads task in humans (Kopp et al., 2016). However, this area has many more functions in cognitive control and decision-making (Domenech & Koechlin, 2015; Ridderinkhof et al., 2004). Teuchies et al. (2016), for example, found evidence that activity in the DLPFC was related to intentional choice in humans. Furthermore, Lin et al. (2020) found that the DLPFC played a role in value computation during decision making in macaque monkeys, and Philiastides et al. (2011) found that the DLPFC is involved in evidence accumulation during perceptual decision making in humans. Quite telling is for example the study by Banca et al. (2016), which looked at evidence accumulation (a measure of impulsivity) on the beads task in both binge drinkers and healthy volunteers, and the relationship with brain volume using voxel-based morphometry. The results showed that less evidence accumulation on the beads task was related to smaller DLPFC and interior parietal volumes (Banca et al., 2016). These results are in line with the findings of Furl and Averbeck (2011), who reported that for take $>$ decline decisions, the size of the BOLD response in the DLPFC showed a significant linear relationship with individual differences in the participants' mean number of samples, with DLPFC responses being greater for participants' who sampled more options.

Another study that investigated evidence accumulation on a type of optimal stopping task is Gluth et al. (2012). In Gluth et al. (2012)'s task, participants were presented with a stock of an unknown value, and they had to decide whether to buy or reject the stock offer. If a participant chose to reject an offer, they were presented with either a positive or a negative rating of the stock (generated randomly), thus providing an indication of its value. There was, however, a fixed monetary cost-to-sample. If a participant chose to buy the stock, they could earn additional bonus payments depending on the value of the stock. There were 32 trials of six stock offers. On this task, activation in the caudate nucleus and anterior insula was associated with evidence accumulation, while activation in the VMPFC, orbitofrontal cortex and ventral striatum was related to the updating of value information (Gluth et al., 2012). Research in the field of neuroeconomics, and specifically on value-based decision-making, extends these findings by identifying that it is the subjective value associated with a stimulus that is encoded in the VMPFC (Brosch & Sander, 2013; Levy & Glimcher, 2012). On the facial attractiveness task (Furl et al., 2019), it is an individual's subjective value of a face that leads to the decision of whether that face will be chosen. As such, the VMPFC is likely to be an important brain area for decision making on the optimal stopping tasks described in this thesis.

To sum up, the prefrontal cortex is involved in many aspects of decision-making on optimal stopping tasks, including, for example, evidence accumulation (Banca et al., 2016; Furl & Averbeck, 2011; Philiastides et al., 2011) and value computation (Gluth et al., 2012; Lin et al., 2020).

1.5.6 Relevance

As outlined in Section 1.2, humans have been observed to display different sampling biases. In the context of optimal stopping problems, the neural correlates of prospective decision-making as listed above have thus far only been described for tasks where humans show an undersampling bias (Costa & Averbeck, 2015; Furl & Averbeck, 2011). Hence, the question remains whether activation in the same areas is correlated with take versus decline decisions on optimal stopping tasks where humans are known to show an oversampling bias. Evidence for this is needed to be able to draw compelling conclusions about the underlying neural correlates of decision making on optimal stopping problems. The functional magnetic resonance imaging (fMRI) study described in Chapter 6 of this thesis addresses this query.

1.6 Research questions and thesis outline

The main aim of this thesis is to understand human sampling biases on full information optimal stopping problems. In other words, to compare human sampling rates to those of an optimal model, and to determine whether humans oversample, undersample, or perhaps sample optimally. Previous literature has found evidence for both undersampling and oversampling biases, which, together with the evidence presented in Sections 1.3, 1.4 and 1.5, led me to my sub-questions:

1. Are sampling biases dependent on the decision-making domain?
2. Can certain task features explain sampling biases?
3. Which brain areas correlate with prospective decision-making on optimal stopping tasks?

1.6.1 Thesis outline

Aside from the introduction, this thesis comprises a methodology chapter, four experimental chapters, and a discussion. A reference list is provided at the end of each chapter. What follows in this section is an outline of the content covered in each chapter. Note that the full rationale for each of my experimental chapters can be found within the chapter itself and will not be presented here.

Chapter 2 focuses on the methodology of this thesis. Within this chapter, I clarify some of the design and methodological decisions I have made. This information is complementary to the methods sections in each individual chapter. Topics covered in Chapter 2 include a critical analysis of the Bayesian ideal observer model, a justification for the choice of online methods, the results of a pilot study, and details regarding my fMRI study.

In Chapter 3 I discuss two studies which together address one of the key questions that emerged from the literature (Section 1.3): are oversampling biases an intrinsic feature of personal mate choice decisions? To answer this question, I investigate sampling decisions when participants choose attractive dates on a client's behalf (Study 1) and when participants evaluate the trustworthiness of faces rather than attractiveness (Study 2). Additionally, I provide a computational model comparison of different models of human behaviour in the Supplementary Materials of Chapter 3.

Following on from the results of Chapter 3, I investigate three distinct image-based decision-making domains in Chapter 4 using different sets of images. Within Chapter 4, I cover two studies aimed at convergent results. The main aim of the chapter is to test whether sampling rates across different decision-making domains depend on a) different over- or undersampling biases, or b) the moments of the generating distributions (as shown for economic-number based tasks, see Section 1.4.1).

As described, the studies reported in Chapters 3 and 4 all employed an image-based optimal stopping task. The three studies discussed in Chapter 5, however, use economic number-based paradigms. The aim of the first two studies described in Chapter 5 is to determine whether the type of stimulus is linked to over- or undersampling biases. The third study takes it one step further and investigates additional task features, as well as the optimal model's generating distribution, and their effects on participants' sampling rate.

The fMRI study detailed in Chapter 6 looks into the neural correlates of evidence seeking, specifically on an optimal stopping task where oversampling biases have been reported in the past (Furl et al., 2019). The aim of this chapter is to determine whether the decision network involved in the decision to stop sampling on economic optimal stopping tasks (where participants sample too few options, see Section 1.5), is also involved in prospective social decisions on a mate choice optimal stopping task.

In the final chapter of this thesis, Chapter 7, I provide a summary of the experimental findings and discuss implications for the field. Furthermore, I provide a critical evaluation of the work, after which I highlight the theoretical contribution of this thesis. I end Chapter 7 with a conclusion, which draws together the research described within the main body of this thesis.

References

Abdelaziz, F. B., & Krichen, S. (2006). Optimal stopping problems by two or more decision makers: A survey. *Computational Management Science*, *4*(2), 89–111. https://doi.org/10.1007/s10287-006-0029-5

Aumann, R. J. (1995). Backward induction and common knowledge of rationality. *Games and Economic Behavior*, *8*(1), 6–19. https://doi.org/10.1016/S0899-8256(05)80015-6

Backwell, P. R. Y., & Passmore, N. I. (1996). Time constraints and multiple choice criteria in the sampling behaviour and mate choice of the fiddler crab, Uca annulipes. *Behavioral Ecology and Sociobiology*, *38*(6), 407–416. https://doi.org/10.1007/s002650050258

Banca, P., Lange, I., Worbe, Y., Howell, N. A., Irvine, M., Harrison, N. A., Moutoussis, M., & Voon, V. (2016). Reflection impulsivity in binge drinking: Behavioural and volumetric correlates. *Addiction Biology*, *21*(2), 504–515. https://doi.org/10.1111/adb.12227

Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23), 12750–12755. https://doi.org/10.1073/pnas.2002312117

Bearden, J. N. (2006). A new secretary problem with rank-based selection and cardinal payoffs. *Journal of Mathematical Psychology*, *50*(1), 58–59. https://doi.org/10.1016/j.jmp.2005.11.003

Bearden, J. N., & Connolly, T. (2007). Multi-attribute sequential search. *Organizational Behavior and Human Decision Processes*, *103*(1), 147–158. https://doi.org/10.1016/j.obhdp.2006.10.006

Bearden, J. N., & Murphy, R. (2007). On Generalized Secretary Problems, In *Uncertainty and Risk: Mental, Formal, Experimental Representations*. Berlin, Heidelberg, Springer. https://doi.org/10.1007/978-3-540-48935-1_11

Bearden, J. N., Rapoport, A., & Murphy, R. (2006). Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Science*, *52*(9), 1437–1449. https://doi.org/10.1287/mnsc.1060.0535

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., Shadlen, M. N., Latham, P. E., & Pouget, A. (2008). Probabilistic Population Codes for Bayesian Decision Making. *Neuron*, *60*(6), 1142–1152. https://doi.org/10.1016/j.neuron.2008.09.021

Beckage, N., Todd, P. M., Penke, L., & Asendorpf, J. (2009). Testing sequential patterns in human mate choice using speed dating, In *Proceedings of the 2009 Cognitive Science Conference*.

Brosch, T., & Sander, D. (2013). Neurocognitive mechanisms underlying value-based decision-making: From core values to economic value. *Frontiers in Human Neuroscience*, *7*.

Campbell, J., & Lee, M. D. (2006). The Effect of Feedback and Financial Reward on Human Performance Solving 'Secretary' Problems. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *28*(28).

Cardinale, E. M., Pagliaccio, D., Swetlitz, C., Grassie, H., Abend, R., Costa, V. D., Averbeck, B. B., Brotman, M. A., Pine, D. S., Leibenluft, E., & Kircanski, K. (2021). Deliberative Choice Strategies in Youths: Relevance to Transdiagnostic Anxiety Symptoms. *Clinical Psychological Science*, 1–11. https://doi.org/10.1177/2167702621991805

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

Chow, Y. S., & Robbins, H. (1963). On optimal stopping rules. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, *2*(1), 33–49. https://doi.org/10.1007/BF00535296

Christian, B., & Griffiths, T. L. (2016). *Algorithms to Live By: The Computer Science of Human Decisions*. Macmillan.

Costa, V. D, & Averbeck, B. B. (2015). Frontal-Parietal and Limbic-Striatal Activity Underlies Information Sampling in the Best Choice Problem. *Cerebral Cortex*, *25*(4), 972–982. https://doi.org/10.1093/cercor/bht286

DeGroot, M. H. (1968). Some Problems of Optimal Stopping. *Journal of the Royal Statistical Society. Series B (Methodological)*, *30*(1), 108–122.

Domenech, P., & Koechlin, E. (2015). Executive control and decision-making in the prefrontal cortex. *Current Opinion in Behavioral Sciences*, *1*, 101–106. https://doi.org/10.1016/j.cobeha.2014.10.007

Eriksson, K., & Strimling, P. (2009). Partner search heuristics in the lab: Stability of matchings under various preference structures. *Adaptive Behavior*, *17*(6), 524–536. https://doi.org/10.1177/1059712309341220

Ferguson, T. S. (1989). Who Solved the Secretary Problem? *Statistical Science*, *4*(3), 282–289. https://doi.org/10.1214/ss/1177012493

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

Fry, H. (2014). The mathematics of love.

Furl, N., & Averbeck, B. B. (2011). Parietal Cortex and Insula Relate to Evidence Seeking Relevant to Reward-Related Decisions. *Journal of Neuroscience*, *31*(48), 17572–17582. https://doi.org/10.1523/JNEUROSCI.4236-11.2011

Furl, N., Averbeck, B. B., & McKay, R. T. (2019). Looking for Mr(s) Right: Decision bias can prevent us from finding the most attractive face. *Cognitive psychology*, *111*, 1–14. https://doi.org/10.1016/j.cogpsych.2019.02.002

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, *61*(313), 35–73. https://doi.org/10.2307/2283044

Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding When to Decide: Time-Variant Sequential Sampling Models Explain the Emergence of Value-Based Decisions in the Human Brain. *Journal of Neuroscience*, *32*(31), 10686–10698. https://doi.org/10.1523/JNEUROSCI.0727-12.2012

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2017). Learning in the Repeated Secretary Problem. *arXiv preprint arXiv:1708.08831*, arxiv 1708.08831.

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning When to Stop Searching. *Management Science*, *66*(3), 1375–1394. https://doi.org/10.1287/mnsc.2018.3245

Griffiths, T. L. (2017). 3 ways to make better decisions - by thinking like a computer.

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339–361. https://doi.org/10.1037/dec0000081

Guan, M., Lee, M. D., & Silva, A. (2014). Threshold Models of Human Decision Making on Optimal Stopping Problems in Different Environments. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*(36).

Haber, S. N., & Knutson, B. (2010). The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology*, *35*(1), 4–26. https://doi.org/10.1038/npp.2009.129

Hauser, T. U., Moutoussis, M., Dayan, P., & Dolan, R. J. (2017). Increased decision thresholds trigger extended information gathering across the compulsivity spectrum. *Translational Psychiatry*, *7*(12), 1–10. https://doi.org/10.1038/s41398-017-0040-3

Hauser, T. U., Moutoussis, M., Purg, N., Dayan, P., & Dolan, R. J. (2018). Beta-Blocker Propranolol Modulates Decision Urgency During Sequential Information Gathering. *The Journal of Neuroscience*, *38*(32), 7170–7178. https://doi.org/10.1523/JNEUROSCI.0192-18.2018

Hill, T. P. (2009). Knowing When to Stop: How to gamble if you must—the mathematics of optimal stopping. *American Scientist*, *97*(2), 126–133.

Hsiao, Y.-C., & Kemp, S. (2020). The effect of incentive structure on search in the secretary problem. *Judgment and Decision Making*, *15*(1), 82–92.

Huettel, S. A., Song, A. W., & McCarthy, G. (2005). Decisions under Uncertainty: Probabilistic Context Influences Activation of Prefrontal and Parietal Cortices. *Journal of Neuroscience*, *25*(13), 3304–3311.

Huq, S. F., Garety, P. A., & Hemsley, D. R. (1988). Probabilistic Judgements in Deluded and Non-Deluded Subjects. *The Quarterly Journal of Experimental Psychology Section A*, *40*(4), 801–812. https://doi.org/10.1080/14640748808402300

Ivy, T. M., & Sakaluk, S. K. (2007). Sequential mate choice in decorated crickets: Females use a fixed internal threshold in pre- and postcopulatory choice. *Animal Behaviour*, *74*(4), 1065–1072. https://doi.org/10.1016/j.anbehav.2007.01.017

Jauhar, S., Fortea, L., Solanes, A., Albajes-Eizagirre, A., McKenna, P. J., & Radua, J. (2021). Brain activations associated with anticipation and delivery of monetary reward: A systematic review and meta-analysis of fMRI studies. *PloS one*, *16*(8). https://doi.org/10.1371/journal.pone.0255292

Kahan, J. P., Rapoport, A., & Jones, L. V. (1967). Decision making in a sequential search task. *Perception & Psychophysics*, *2*(8), 374–376. https://doi.org/10.3758/BF03210074

Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of Visual Attention in the Human Cortex. *Annual Review of Neuroscience*, *23*(1), 315–341. https://doi.org/10.1146/annurev.neuro.23.1.315

Kennerley, S. W., Walton, M. E., Behrens, T. E. J., Buckley, M. J., & Rushworth, M. F. S. (2006). Optimal decision making and the anterior cingulate cortex. *Nature Neuroscience*, *9*(7), 940–947. https://doi.org/10.1038/nn1724

Kiani, R., & Shadlen, M. N. (2009). Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science*, *324*(5928), 759–764. https://doi.org/10.1126/science.1169405

Klein, T. A., Ullsperger, M. M. P., & Danielmeier, C. (2013). Error awareness and the insula: Links to neurological and psychiatric diseases. *Frontiers in Human Neuroscience*, *7*, 1–14. https://doi.org/10.3389/fnhum.2013.00014

Knutson, B., Adams, C. M., Fong, G. W., & Hommer, D. (2001a). Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens. *Journal of Neuroscience*, *21*(16). https://doi.org/10.1523/JNEUROSCI.21-16-j0002.2001

Knutson, B., Fong, G. W., Adams, C. M., Varner, J. L., & Hommer, D. (2001b). Dissociation of reward anticipation and outcome with event-related fMRI. *Neuroreport*, *12*(17), 3683–3687.

Knutson, B., Rick, S., Wimmer, G. E., Prelec, D., & Loewenstein, G. (2007). Neural Predictors of Purchases. *Neuron*, *53*(1), 147–156. https://doi.org/10.1016/j.neuron.2006.11.010

Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, *14*(3), 381–392. https://doi.org/10.1016/0167-2681(90)90065-L

Kolling, N., Behrens, T. E. J., Wittmann, M. K., & Rushworth, M. F. S. (2016a). Multiple signals in anterior cingulate cortex. *Current Opinion in Neurobiology*, *37*, 36–43. https://doi.org/10.1016/j.conb.2015.12.007

Kolling, N., Scholl, J., Chekroud, A., Trier, H. A., & Rushworth, M. F. S. (2018). Prospection, Perseverance, and Insight in Sequential Behavior. *Neuron*, *99*(5), 1069–1082. https://doi.org/10.1016/j.neuron.2018.08.018

Kolling, N., Wittmann, M. K., Behrens, T. E. J., Boorman, E. D., Mars, R. B., & Rushworth, M. F. S. (2016b). Value, search, persistence and model updating in anterior cingulate cortex. *Nature Neuroscience*, *19*(10), 1280–1285. https://doi.org/10.1038/nn.4382

Kopp, B., Seer, C., Lange, F., Kluytmans, A., Kolossa, A., Fingscheidt, T., & Hoijtink, H. (2016). P300 amplitude variations, prior probabilities, and likelihoods: A Bayesian ERP study. *Cognitive, Affective, & Behavioral Neuroscience*, *16*(5), 911–928. https://doi.org/10.3758/s13415-016-0442-3

Kuhnen, C. M., & Knutson, B. (2005). The Neural Basis of Financial Risk Taking. *Neuron*, *47*(5), 763–770. https://doi.org/10.1016/j.neuron.2005.08.008

Lee, M. D. (2006). A Hierarchical Bayesian Model of Human Decision-Making on an Optimal Stopping Problem. *Cognitive Science*, *30*(3), 1–26. https://doi.org/10.1207/s15516709cog0000_69

Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, *22*(6), 1027–1038. https://doi.org/10.1016/j.conb.2012.06.001

Lin, Z., Nie, C., Zhang, Y., Chen, Y., & Yang, T. (2020). Evidence accumulation for value computation in the prefrontal cortex during decision making. *Proceedings of the National Academy of Sciences*, *117*(48), 30728–30737. https://doi.org/10.1073/pnas.2019077117

Loued-Khenissi, L., Pfeuffer, A., Einhäuser, W., & Preuschoff, K. (2020). Anterior insula reflects surprise in value-based decision-making and perception. *Neuroimage*, *210*, 116549. https://doi.org/10.1016/j.neuroimage.2020.116549

Mantonakis, A., Rodero, P., Lesschaeve, I., & Hastie, R. (2009). Order in Choice: Effects of Serial Position on Preferences. *Psychological Science*, *20*(11), 1309–1312. https://doi.org/10.1111/j.1467-9280.2009.02453.x

Milinski, M., & Bakker, T. (1992). Costs influences sequential mate choice in sticklebacks, gasterosteus aculeatus. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *250*(1329), 229–233. https://doi.org/10.1098/rspb.1992.0153

Miller, G. F., & Todd, P. M. (1998). Mate choice turns cognitive. *Trends in Cognitive Sciences*, *2*(5), 190–198. https://doi.org/10.1016/S1364-6613(98)01169-3

Palley, A. B., & Kremer, M. (2014). Sequential Search and Learning from Rank Feedback: Theory and Experimental Evidence. *Management Science*, *60*(10), 2525–2542. https://doi.org/10.1287/mnsc.2014.1902

Papadimitriou, C. H., & Tsitsiklis, J. N. (1987). The Complexity of Markov Decision Processes. *Mathematics of Operations Research*, *12*(3), 441–450. https://doi.org/10.1287/moor.12.3.441

Pesaran, B., Nelson, M. J., & Andersen, R. A. (2008). Free choice activates a decision circuit between frontal and parietal cortex. *Nature*, *453*(7193), 406–409. https://doi.org/10.1038/nature06849

Petruccelli, J. D. (1980). On a Best Choice Problem with Partial Information. *Annals of Statistics*, *8*(5), 1171–1174. https://doi.org/10.1214/aos/1176345156

Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., & Blankenburg, F. (2011). Causal Role of Dorsolateral Prefrontal Cortex in Human Perceptual Decision Making. *Current Biology*, *21*(11), 980–983. https://doi.org/10.1016/j.cub.2011.04.034

Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233–238. https://doi.org/10.1038/22268

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

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The Role of the Medial Frontal Cortex in Cognitive Control. *Science*, *306*(5695), 443–447. https://doi.org/10.1126/science.1100301

Sakaguchi, M. (1961). Dynamic programming of some sequential sampling design. *Journal of Mathematical Analysis and Applications*, *2*(3), 446–466. https://doi.org/10.1016/0022-247X(61)90023-3

Samuels, S. M. (2004). Why do these quite different best-choice problems have the same solutions? *Advances in Applied Probability*, *36*(2), 398–416. https://doi.org/10.1239/aap/1086957578

Schuck, N. W., Gaschler, R., Wenke, D., Heinzle, J., Frensch, P. A., Haynes, J.-D., & Reverberi, C. (2015). Medial Prefrontal Cortex Predicts Internally Driven Strategy Shifts. *Neuron*, *86*(1), 331–340. https://doi.org/10.1016/j.neuron.2015.03.015

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

Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, *13*(4), 391–411. https://doi.org/10.1002/1099-0771(200010/12)13:4<391::AID-BDM359>3.0.CO;2-I

Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: Seeing and deciding. *Proceedings of the National Academy of Sciences*, *93*(2), 628–633. https://doi.org/10.1073/pnas.93.2.628

Shepp, L. A. (1969). Explicit Solutions to Some Problems of Optimal Stopping. *The Annals of Mathematical Statistics*, *40*(3), 993–1010. https://doi.org/10.1214/aoms/1177697604

Shu, S. B. (2008). Future-biased search: The quest for the ideal. *Journal of Behavioral Decision Making*, *21*(4), 352–377. https://doi.org/10.1002/bdm.593

Siegmund, D. O. (1967). Some Problems in the Theory of Optimal Stopping Rules. *The Annals of Mathematical Statistics*, *38*(6), 1627–1640.

Stewart, T. J. (1978). Optimal Selection from a Random Sequence with Learning of the Underlying Distribution. *Journal of the American Statistical Association*, *73*(364), 775–780. https://doi.org/10.1080/01621459.1978.10480098

Teuchies, M., Demanet, J., Sidarus, N., Haggard, P., Stevens, M. A., & Brass, M. (2016). Influences of unconscious priming on voluntary actions: Role of the rostral cingulate zone. *NeuroImage*, *135*, 243–252. https://doi.org/10.1016/j.neuroimage.2016.04.036

Thielscher, A., & Pessoa, L. (2007). Neural Correlates of Perceptual Choice and Decision Making during Fear–Disgust Discrimination. *Journal of Neuroscience*, *27*(11), 2908–2917.

Van der Leer, L., Hartig, B., Goldmanis, M., & McKay, R. (2015). Delusion proneness and 'jumping to conclusions': Relative and absolute effects. *Psychological Medicine*, *45*(6), 1253–1262. https://doi.org/10.1017/S0033291714002359

Wittmann, M. K., Kolling, N., Akaishi, R., Chau, B. K. H., Brown, J. W., Nelissen, N., & Rushworth, M. F. S. (2016). Predictive decision making driven by multiple time-linked reward representations in the anterior cingulate cortex. *Nature Communications*, *7*(1). https://doi.org/10.1038/ncomms12327

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer Sequential Search: Not Enough or Too Much? *Marketing Science*, *22*(4), 503–519. https://doi.org/10.1287/mksc.22.4.503.24909

Chapter 2

Methodology

Within this chapter, I summarise the key methodology used within this thesis to address my main aim of explaining human oversampling biases on full information optimal stopping problems. As discussed in Chapter 1, the rationale for this research aim is the oversampling bias reported by Furl et al. (2019), which was in contrast with the majority of previous research on full information problems, which reported undersampling instead (e.g., Baumann et al., 2020; Cardinale et al., 2021; Costa & Averbeck, 2015). To identify human biases, participants' behaviour must be compared to a model of optimality. For the studies included in this thesis, I have chosen to compare participants' sampling behaviour to that of a Bayesian ideal observer model. In Section 2.1, I present a critical evaluation of this model as well as alternative approaches, thereby justifying my choice for the Bayesian model. A second choice I have made is to conduct some studies in the lab and some online. In Section 2.2, I explain why this thesis does not only include traditional lab-based behavioural studies, but includes online studies as well that utilise novel technologies. Also described in this section are the results of a pilot study which I conducted to determine the reliability of using online methods for full information optimal stopping problems. I note that a complete description of the methodology used for each of my studies can be found in the methods section of the chapter concerned.

2.1 Bayesian models versus heuristics models

In this section, I provide a critical evaluation of the Bayesian ideal observer model to which I compare participants' behaviour in my studies, showing why it is the most suitable for answering my research questions. First, I will briefly describe the assumptions of Bayesian modelling in general, after which I focus specifically on my model specification. Finally, I will discuss an alternative non-optimal framework of heuristics models.

2.1.1 Bayesian modelling

Bayesian models are often used for developing computational models of human behaviour (Chambers & Kording, 2018). The foundation for Bayesian models lies in probability theory and particularly Bayes' theorem, which describes the probability of an outcome, based on prior knowledge and beliefs (hypothesis h). Bayes' theorem then updates its prior to posterior probabilities, based on the data encountered (d) (Tenenbaum et al., 2011). Mathematically, Bayes' theorem is stated as

$$P(h \mid d) = \frac{P(d \mid h)P(h)}{P(d)} \tag{2.1}$$

with $P(h \mid d)$ being the posterior probability, $P(h)$ and $P(d)$ being the prior probabilities of h and d, and $P(d \mid h)$ being the likelihood of h given d.

Throughout my thesis, I employ a model of optimality know as the *Bayesian ideal observer model*. The term ideal observer refers to what is known in the literature as the *mathematical solution*. The mathematical solution stands for the characterisation of the problem people are solving and its ideal solution (Griffiths et al., 2012; Marr, 1982) (i.e., Equation 2.1). This means that the Bayesian ideal observer model does not imply that people are actually computing the optimal solution through Bayes' theorem (Griffiths et al., 2012). Instead, it simply characterises the problem people are solving, rather than the mechanism by which they might be solving it (Griffiths et al., 2012).

In the Supplementary Materials of Chapter 3, I compare participants' behaviour not only to the Bayesian ideal observer model, but to three additional Bayesian computational models of human behaviour (see Sidebar 4). These computation models appeal to the so-called *algorithmic procedure* (Griffiths et al., 2012; Marr, 1982; Shi et al., 2010), which refers to the cognitive processes underlying decision-making (Griffiths et al., 2012). In other words, the computational models specify the computational mechanisms that humans might be using to solve a task. These models of decision-making biases can be viewed as hypotheses about cognitive processes. As Al-Shawaf and Buss (2011) state, Bayesian models could assist in the endeavor to elucidate computational mechanisms underlying human behaviour. The benefit of using Bayesian models for the purpose of identifying the cognitive processes that may underlie sampling biases on optimal stopping tasks is that they clarify the bias that is being approximated and the nature of the bias approximation (Griffiths et al., 2012). In my case, I parameterise the optimal model to formulate hypotheses about where in its otherwise optimal workflow this parameterised bias might exert its influence. Another advantage of using Bayesian models is that Bayesian models allow for the explicit integration of a prior hypothesis h with new data d (Heit & Erickson, 2011). This incorporation of a prior hypothesis resembles a crucial element of human decision making, which is often based on prior knowledge and experience (Heit & Erickson, 2011).

Sidebar 4: Bayesian theoretical models. The Bayesian ideal observer model can be adapted to become a theoretical model of human behaviour instead of an optimal model by altering the input and/or model values. In this way, the cognitive processes underlying sampling biases can be approximated. Consider, for example, the cost-to-sample parameter. By treating the cost-to-sample parameter as a free parameter, it can be optimised to resemble human sampling behaviour. Optimisation of free parameters can be achieved in several ways. For example, one could select a value for the parameter that produces sampling rates equal to that of the mean participant (Furl et al., 2019), or one could employ methods used in the machine learning literature and minimise the negative log likelihood. The result is a theoretical model with a cost-to-sample parameter that is either positive or negative. If after optimisation the cost-to-sample parameter is positive, it could be the case that participants make decisions with an intrinsic time cost in mind which leads to undersampling (see Section 1.4.2). On the other hand, if after optimisation the cost-to-sample parameter is negative, participants could find sampling intrinsically rewarding, and consequently oversample. I consider this and other such parameterisations in the Supplementary Materials of Chapter 3.

There is an ongoing debate in the literature regarding the contribution of Bayesian models to cognitive theory (e.g., Al-Shawaf & Buss, 2011; Bowers & Davis, 2012; Griffiths et al., 2012; Heit & Erickson, 2011; Jones

& Love, 2011). Jones and Love (2011) argue that before Bayesian models can be used to understand the cognitive processes underlying human behaviour, more clarity is needed regarding what can be claimed from the findings. In other words, it should be made clear whether the Bayesian model appeals to the mathematical solution or the algorithmic procedure. As mentioned above, the Bayesian computational models described in the Supplementary Materials of Chapter 3 appeal to the algorithmic procedure, while the Bayesian ideal observer model, as implemented everywhere else throughout this thesis, appeals to the mathematical procedure.

2.1.2 Bayesian ideal observer model

As discussed in Chapter 1, some of the first to propose an optimal solution to full information optimal stopping problems are Gilbert and Mosteller (1966). Since in full information problems the actual value of an option is presented, rather than its relative rank, a decision maker can decide on a threshold for each position in the sequence at which they will choose an option. When a value exceeds the threshold for that position in the sequence, it is chosen by the decision maker (Lee et al., 2005). The Bayesian ideal observer model is the same as Gilbert and Mosteller (1966) in that it computes the expected values of future options under the assumed-to-be normal generating distribution [1] via a backwards induction technique (see Sidebar 1). The original model by Gilbert and Mosteller (1966) knows the mean and variance of the generating distribution for all sequence options, and that the reward values of the sequence ranks are equal to the option values. The Bayesian ideal observer model extends this framework by adding to it 1) a generating distribution that is initialised with a prior distribution which is then updated after each new sample using Bayes' rule, 2) a cost-to-sample parameter, and 3) functionality for the researcher to apply any arbitrary reward function to the choice outcomes. Theoretically, at each position in the sequence, the ideal observer model computes the respective values for choosing the option and declining the option, and chooses the one with the highest value. The value for declining an option can be considered the choice threshold, as no option is chosen unless the value for choosing an option exceeds the value for declining an option. The choice threshold is dynamic, and can change depending on the position in the sequence.

[1] Recall that the generating distribution refers to the distribution that generates the values of the options in a sequence.

The Gilbert and Mosteller model (Gilbert & Mosteller, 1966) has historically been regarded in the optimal stopping literature as the optimal solution to full information problems (e.g., Baumann et al., 2020; Guan et al., 2015; Lee et al., 2005; Samuel-Cahn, 1996). The Bayesian ideal observer model, which is an advanced version of the Gilbert and Mosteller model, has also gained some attention in recent years, with previous research employing the model on full information problems including number-based tasks (Cardinale et al., 2021; Costa & Averbeck, 2015), the facial attractiveness task (Furl et al., 2019), and the beads task (Furl & Averbeck, 2011). I will now outline my reasons for choosing the Bayesian ideal observer model as the model of optimality for the studies within this thesis.

Bayesian approaches have been successful in modelling human behaviour in a plethora of complex domains including, for example, language (Griffiths et al., 2007), reasoning (Kemp & Tenenbaum, 2009), categorisation (Anderson, 1991), and magnitude estimation (Petzschner et al., 2015). Like the Bayesian ideal observer model, most of these Bayesian cognitive models appeal to the mathematical solution as described in Section 2.1 (Sanborn & Chater, 2016). This shows that the use of Bayesian (optimal) models finds wide support within the behavioural science literature, even extending to other fields including neuroscience, economics and philosophy (Chater et al., 2020). Because of this wide implementation of Bayesian approaches, the theoretical contributions of this thesis could potentially apply to a wider range of domains (Griffiths et al., 2012).

Furthermore, the facial attractiveness task and its irregular findings of oversampling form the outset of this thesis. To further investigate the oversampling bias reported by Furl et al. (2019), I chose to use the same model of optimality as these authors. This means that not only do I use the Bayesian ideal observer model, the model parameters (i.e., the mean and variance of the prior distribution, the reward function and the cost-to-sample parameter) were also specified in the same way as Furl et al. (2019). As mentioned before, the Bayesian ideal observer model is an advanced version of the Gilbert and Mosteller model (Gilbert & Mosteller, 1966) which is historically used in the literature for full information problems. The added functionalities, particularly the implementation of Bayes' rule and the reward function, are suited to model the task features as implemented in the facial attractiveness task, as Furl et al. (2019) have demonstrated by using the Bayesian ideal observer model for their studies.

Having provided a justification for my choice for the Bayesian ideal observer model as the model of optimality in my studies, I will now discuss a set of alternative models of human behaviour which can be characterised as heuristics models and explain why they are not suited to answering my research questions.

2.1.3 Heuristics models

As discussed in Chapter 1, the optimal stopping problem originated as a mathematics problem. Optimal stopping rules, such as the 37% rule for the classical secretary problem, are therefore mathematically optimal solutions to the problem. The development that followed was that researchers started comparing human behaviour to these models of optimality. When they found that humans did not (always) follow the optimal rule, they started describing human sampling on optimal stopping tasks in terms of heuristics models. Heuristics are matched to a particular environment and describe a decision mechanism that people can use to make good decisions within internal (cognitive) and external (environmental) bounds (Todd & Gigerenzer, 2003). They do *not* represent the optimal solution to an optimal stopping problem. In fact, some researchers that focus on heuristics models appear to reject the usefulness of models of optimality due to their psychological implausibility (Todd & Gigerenzer, 2003). Within this section I discuss some of the benefits and limitation of heuristics models and demonstrate why heuristics models are unsuitable for answering my research questions.

Todd and Gigerenzer (2001) argue that it is unnecessary for researchers to compute the optimal solution and to use it to compare human behaviour on optimal stopping tasks, because heuristics are able to describe human decision making on these tasks. It is true that experimental studies have found evidence for heuristics models fitting well with participants' sampling behaviour on, for example, the classical secretary problem (Seale & Rapoport, 1997) and the secretary problem with cardinal payoffs (Todd & Miller, 1999). However, these rules can quickly break down when the environment, i.e., the task design, is changed (Todd & Gigerenzer, 2003). The fact that heuristics models require ad hoc specification of heuristic devices (Lee, 2006) has led to a wide diversity in heuristics models in the optimal stopping literature. I illustrate this point here by describing several heuristics that have been proposed for mate choice paradigms (e.g., paradigms like Furl et al., 2019, who investigated human sampling biases on the full information facial attractiveness task).

- **Best-of-N rule**: sampling a certain number of options, after which the best of those seen is chosen (Todd & Miller, 1999; Valone et al., 1996). This heuristic is limited to paradigms where recall is permitted, and has been found to explain not only human mate choice behaviour, but that of animals as well (Seeley & Buhrman, 2001).

- **Take-the-next-best rule**: refrain from making any decisions for a fixed proportion of options in the sequence, after the next top-ranked option should be chosen (Beckage et al., 2009; Todd & Miller, 1999). In the context of the classical secretary problem this heuristic is also known as the cutoff model (Dudey & Todd, 2001; Sang et al., 2020; Seale & Rapoport, 1997, 2000).

- **Reservation rule**: stop sampling if an option is presented that is of sufficiently high value, i.e., exceeds a pre-set threshold or aspiration level (Hey, 1987; Todd & Miller, 1999).

- **Threshold rule**: comprises a number of heuristics including the take-the-next-best rule and reservation rule. An option has to exceed a certain value threshold for it to be chosen (Guan et al., 2014; Todd & Miller, 1999). Variations of this heuristic exist that describe situations where humans change their threshold over time, for example as a result of feedback (Beckage et al., 2009; Todd & Miller, 1999) or because they are approaching the end of the sequence (the 'prettier-at-closing-time effect', Eriksson & Strimling, 2009).

The list above illustrates both the diversity of heuristics models in the optimal stopping literature, as well as the dependency of some of these heuristics on specific task features. For example, Lee et al. (2005) reported that the cutoff model, which was found to perform well on classical secretary problems, was not a suitable model of human decision making on a full information problem. Instead, on full information problems (non)linear threshold models have been found to capture human sequential decision-making (Baumann et al., 2020; Thomas et al., 2021). However, like most heuristics models, these (non)linear threshold models make a number of strong theoretical assumptions including that "a participant uses the same thresholds for each problem, and there is no learning, adaptation, or self-regulation of the thresholds after the practice problems have been completed." (p. 347, Guan & Lee, 2018). Therefore, while heuristics can provide accurate ecologically-valid descriptions of human behaviour, their application is limited. Bayesian models, on the other hand, can be used as powerful, robust yet flexible tools for modelling human behaviour. As Griffiths et al. (2012) points out, Bayesian models can be modified to represent hypotheses about the cognitive process that underlie human behaviour. The benefit of using Bayesian models to describe human behaviour is that they are less task-dependent than heuristics models because they stem directly from the optimal solution, rather than the ad hoc solutions implemented in heuristics models. As a proof of concept, in the Supplementary Materials of Chapter 3 I demonstrate that Bayesian computational models of human behaviour better fit participants' data than a heuristics model.

To summarise, I have claimed above that Bayesian models show the most potential, as heuristics models do not generalise well to problems outside their narrow domains and may therefore not be able to provide the standardisation needed to compare the optimal stopping literature. As outlined in the Introduction, sampling behaviour on optimal stopping tasks is heavily task-dependent. Without a standard measure (i.e., an optimality benchmark) against which to compare human sampling rates, it becomes nearly impossible to explain differences in sampling biases, and consequently, nearly impossible to explain human oversampling biases on full information problems. Heuristics models could potentially be used to describe the cognitive processes underlying sampling biases, but the Bayesian ideal observer model can be modified to achieve the same goal. This thesis preeminently focuses on comparisons between human behaviour and optimality benchmarks, for which heuristics are by definition sub-optimal. However, one sub-element of this thesis (Supplementary Materials of Chapter 3) does compare Bayesian theoretical models to human behaviour. In this model comparison, I also compare these models against the most prominent heuristic for modelling optimal stopping problems; the cutoff model (Dudey & Todd, 2001; Sang et al., 2020; Seale & Rapoport, 1997, 2000).

2.2 Online methods for behavioural studies

Some behavioural studies described in this thesis follow a more traditional lab-based protocol, while others are designed and executed using online methods. Within this section, I will provide further clarification for why I chose to use online methods, how I mitigated some of the concerns regarding online methods, and I report the results of a pilot study.

2.2.1 Benefits

Some of the benefits of online methods include ease-of-use, low cost, and quick collection of large, representative samples (Anwyl-Irvine et al., 2020b; Casler et al., 2013; Craig et al., 2013; Hauser et al., 2018; Palan & Schitter, 2018; Sheehan, 2018). The latter especially has been identified as a desideratum. Peterson (2001) for example, concluded from their large second-order meta-analysis that the effects observed in a typical student population cannot simply be extended to a non-student (adult) population. Therefore, all online studies described in this thesis include a random sample of participants, consisting of both students and non-students.

Additionally, online methods have been identified as a promising tool to improve the replicability of scientific research. A variety of traditional laboratory findings have been replicated in online versions, ranging from studies on cognition and cognitive psychology (Carpenter et al., 2019; Crump et al., 2013; Kochari, 2019) to judgement and decision-making (Paolacci et al., 2010) and learning (Casler et al., 2013).

There are myriads of online recruitment platforms available, with the two biggest and most popular being Amazon Mechanical Turk (MTurk) and Prolific. For my behavioural studies, I have chosen to use Prolific, as this service has been found to have high internal reliability, high transparency (towards both participants and researchers), good reproducibility, and a lower level of dishonesty, compared to other platforms (Palan & Schitter, 2018; Peer et al., 2017).

2.2.2 Concerns

Some researchers have raised concerns regarding the use of online recruitment platforms. Ford (2017), for example, noted that a sample obtained through the MTurk platform may contain 'speeders' and 'cheaters'. Speeders are participants who do not pay attention to the task and whose primary goal is to finish as quickly as possible. Cheaters are participants who provide incorrect information about themselves in order to be included in more studies so they can make more money (Ford, 2017). Kees et al. (2017b) note that the concerns raised by Ford (2017) are applicable to most online data sources. To mitigate these concerns and to improve data quality, researchers have sought to provide guidelines for online data collection (Hauser et al., 2018; Kees et al., 2017a; Sheehan, 2018). Here, I summarise some of these guidelines including how I implement them in my online behavioural studies.

- **Ensuring attentiveness** (Hauser et al., 2018; Kees et al., 2017a; Sheehan, 2018): including attention checks to ensure attentiveness is generally considered good practice, and it is encouraged by Prolific. I included attention checks in the majority of my studies, unless the study was very short ($\sim$10min or less), in which case including attention checks is deemed unnecessary. Participants who do not pass $> 75\%$ of the attention checks are excluded from the data analysis.

- **Language comprehension** (Hauser et al., 2018): as participants may be recruited from anywhere in the world, it is important to set certain prerequisites to ensure language comprehension. Hauser et al. (2018) suggest to include "only those in the geographic regions that are likely to speak the language of the survey fluently" (p. 18). I implement this suggestion in the online study described in Chapter 4. For the studies described in Chapter 5, I use one of Prolific's pre-screening questions to only include participants who had indicated they were fluent in the English language.

- **Minimizing non-naivete** (Hauser et al., 2018): there is some evidence that repeated exposure to optimal stopping problems can lead to learning and improved task performance (Goldstein et al., 2017). Therefore, following the recommendations made by Hauser et al. (2018), I have made every possible attempt to prevent participants from completing more than one of my studies, including the use of Prolific's exclusion criteria.

In the context of this thesis, one limitation of using online methods is that online studies are more limited in the duration that participants can be asked to be attentive for. Therefore, in this thesis, online studies will be reserved for assessing sampling bias only. That is, I will compare participants' sampling rates from online studies only to the Bayesian ideal observer model described above (Section 2.1.2). Theoretical modelling of the cognitive processes underlying decision-making, as implemented in Study 2 of Furl et al. (2019), involves a model fitting procedure that is likely to be more effective with more trials per participant. Because it takes at least an hour to obtain the amount of data needed to facilitate model fitting, which exceeds the duration of an ordinary online study, the computational model comparisons as implemented in the Supplementary Materials of Chapter 3 (which require model fitting) are performed on longer, lab-based studies.

2.2.3 Pilot study

I conducted a pilot study to determine whether online methods can reliably be used to investigate sampling biases on optimal stopping problems. I chose to pre-test the facial attractiveness paradigm (described in Chapter 1, Section 1.3) as this paradigm is of particular importance to answering my research questions, and it allows for a direct comparison with previous research (Furl et al., 2019). If the results replicate the oversampling bias reported in lab studies (Furl et al., 2019), the pilot study is considered successful. Additionally, lab studies found that participants on average ended up with lower-ranked options compared to the optimal model, which I investigate here as well.

Participants

The pilot study was approved by Royal Holloway, University of London's ethics board. Informed consent was obtained from all participants before the start of the study, in accordance with the Declaration of Helsinki. The sample size (N = 20) was based on power analyses derived from Furl et al. (2019). Gorilla Experiment Builder (Anwyl-Irvine et al., 2020a) was used to create and host the pilot study, and participants were recruited through the online recruitment service Prolific (Prolific, 2014). Two participant prerequisites were set, the first being age (between 18 and 35), as this roughly matched the age range of the faces shown in the study. The second prerequisite was nationality (either United Kingdom (UK), Ireland, United States (US), Canada, Australia or New Zealand), which was set under the assumption that participants with these nationalities would have a good command of the English language, and would therefore be able to sufficiently understand the instructions and informed consent form. One participant exceeded the time limit set on the Gorilla (2x estimated completion time), which is why data from only 19 participants was saved on Gorilla's server (14 female participants, 15 UK nationals, M_{age} = 25.0 years, SD_{age} = 3.25 years).

Materials and methods

The methods for the pilot study were broadly taken from the methods used across the three studies described in Furl et al. (2019). After presenting the instruction sheet and consent form, participants were asked to choose whether they would like to rate (and date) males or females. Based on their answers, each was shown either male or female faces throughout the study.

In phase one of the study, participants rated 180 faces in total (90 unique faces, all rated twice) using a slider scale ranging from very unattractive (value 1) to very attractive (value 100), with the initial marker position set to the middle of the slider. Final attractiveness ratings were computed from the mean of the two ratings. Face images were randomly selected from a larger set of 426 images as used by Furl et al. (2019). An attention check[2] was added to phase one to compensate for the unsupervised nature of online data collection, which was passed by all participants.

In the second phase of the study, participants were shown six sequences of eight faces each. Faces were randomly sampled from the entire distribution of faces that had been rated in phase one. Participants attempted to choose the most attractive person in the sequence as their date, with the restriction that they could not return to a previously rejected date. The number of options remaining was shown at the top of the screen, and the rejected 'dates' were shown at the bottom of the screen. When a participant made a choice, they had to advance through a series of grey squares that replaced the remaining faces. This ensured that participants could not finish the study early by choosing an early option. The entire study was self-paced - participants advanced by using their mouse to click on the buttons on the screen. If the last face in the sequence was reached, that person became their 'date' by default. After finishing a sequence, participants were directed to a feedback screen displaying the participant's chosen face and the text: "Here is your

[2]For a full description of the attention check see the Supplementary Materials of either Chapter 4 (Section S1), or Chapter 5 (Text A).

new date! How rewarding is your choice?". Like Furl et al. (2019), participants rated the reward value of their choice on a 9-point scale. Responses to the feedback screen were analysed as a sanity check, but were otherwise irrelevant to the goal of the pilot study.

The two dependent variables of interest are the position of the chosen image in the sequence (i.e., number of samples), and the rank of the chosen image (out of the images in the sequence). Both variables are a mean value over six sequences for each participant, and can be compared to the results of Furl et al. (2019).

Data analysis

Participants' sampling behaviour was compared to the Bayesian ideal observer model (Furl et al., 2019), where performance is Bayesian optimal and the cost-to-sample parameter was fixed to zero (see Section 2.1.2). For a mathematical description of the model, see, for example, Chapter 4, Section 2.2.3. Like the model implementation used in Furl et al. (2019), the Bayesian ideal observer model received as input for each participant the subjective values (final attractiveness ratings) of the sequence options as presented to the participant in phase two. To approximate normality, ratings were log transformed for each participant before being put into the model. The comparison of participants' sampling behaviour to the ideal observer model was done using MATLAB version 2015b (MATLAB, 2015). A p value of $< .05$ was considered significant.

Results

Using standardised Cronbach's alpha, I found a good internal consistency between participants' mean attractiveness rating and the self-reported reward value of a chosen face (α = .854). This supports the validity of the experimental paradigm and indicates that participants' subjective ratings correspond well to endogenous reward values. Furthermore, as a sanity check, I examined whether there was a relationship between participants' self-reported reward value and the position of the chosen face in the sequence. Expected is a negative correlation between reward values and sequence position, as faces at the last position in the sequence get chosen by default, whilst faces at earlier positions get chosen because they are (supposedly) sufficiently attractive. Indeed, using Spearman's rho, I found a negative correlation between the self-reported reward value and the position of the chosen face in the sequence (r_s = -.74, $p < .001$).

The main results of the pilot study indicated that, compared to a Bayesian ideal observer model, participants oversampled ($t(18) = 5.9$, $p < .001$) and ended up with lower-ranked faces ($t(18) = -2.7$, $p = .014$) (Figure 2.1).

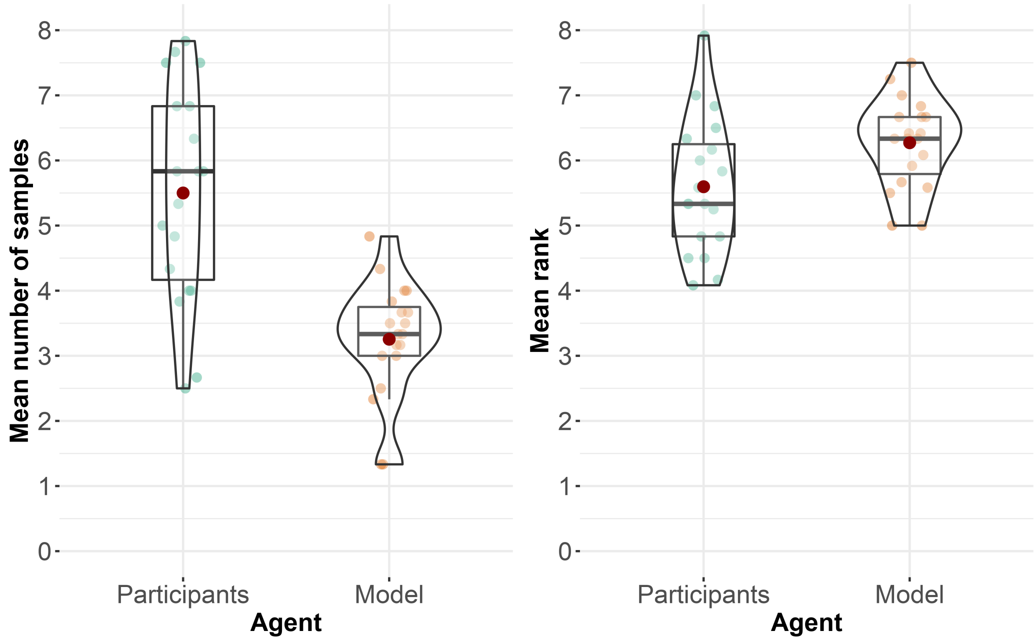

FIGURE 2.1: Distributions of the mean number of samples and the mean rank for participants versus the Bayesian ideal observer model. The red dots represent the mean, horizontal black lines represent the median, boxes show the 25% and 75% quantiles, and the whiskers represent the 95% confidence intervals.

Conclusion

With a sample size of only 19 participants, the pilot study replicated the findings of previous lab-based research (Furl et al., 2019). This highlights the robustness of the facial attractiveness paradigm, and indicates that online methods (specifically Gorilla and Prolific) may provide a reliable alternative to traditional methods in optimal stopping research. Furthermore, I interpreted the results of the sanity check as indicating that participants behaved rationally and in line with the expectations of how they should interact with the task. The good internal consistency between participants' mean attractiveness rating and the self-reported reward value of a chosen face is suggestive that there are no changes in subjective ratings before and after completing the optimal stopping task. This is consistent with previous research that has varied the number of attractiveness ratings before the optimal stopping task, but found no differences in participants' sampling behaviours (three ratings in Studies 1 and 2 and two ratings in Study 3; Furl et al., 2019. However, my analysis of the self-reported reward value of the chosen face might be biased by the fact that only self-reported reward values are available for options that were chosen by the participant, and not available for the other competing options in the sequences. Additionally, the question to capture the participants' ratings was phrased differently before and after the optimal stopping task ("Rate this face on its attractiveness" versus "How rewarding is your choice"). The high reliability of different ratings of the same face within phase 1, the replicated oversampling despite different numbers of phase 1 ratings per face in different studies, and the concordance of phase 1 ratings with self-reported reward values of chosen faces in this pilot study are all suggestive that perceived attractiveness of a given face is stable over time. However, a more direct and less biased approach to explore whether individual differences in rating (in)consistency relate to sampling behaviour, which future research could address, might be to repeat the rating phase completely after the optimal stopping task.

Of note is that I set the time limit for my pilot study to twice the estimated completion time, following the rule of thumb proposed by Sheehan (2018) (2-3 times the length of time it should take to complete the study). However, this time limit was found to be too strict as data from one participant was not saved on Gorilla's server as a result of this. Hence, for all studies described in the experimental chapters, I utilise Prolific's automated 'Maximum Time Allowed' feature which is based on the estimated study duration set by the researcher. This successfully prevented loss of data as observed for the pilot study. Additionally, I have addressed concerns regarding online data collection (Hauser et al., 2018; Kees et al., 2017a; Sheehan, 2018) in my pilot study by implementing an attention check and ensuring language comprehension. Similar methods to those described above have been applied in Chapters 4 and 5 where I implement online methods.

2.3 Summary

I have now laid out the rationale regarding the key methodology employed within this thesis, including a critical analysis of those methods. I have compared my choice for the Bayesian ideal observer model to alternative heuristics models and I have detailed the considerations made regarding the use of online methods including proof of concept in the form of a pilot study. What follows next are four experimental chapters which employ the methodology outlined in this chapter.

References

Al-Shawaf, L., & Buss, D. (2011). Evolutionary psychology and Bayesian modeling. *Behavioral and Brain Sciences*, *34*(4), 188–189. https://doi.org/10.1017/S0140525X11000173

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409–429. https://doi.org/10.1037/0033-295X.98.3.409

Anwyl-Irvine, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2020a). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods*, *52*(1), 388–407. https://doi.org/10.3758/s13428-019-01237-x

Anwyl-Irvine, A. L., Dalmaijer, E. S., Hodges, N., & Evershed, J. K. (2020b). Online Timing Accuracy and Precision: A comparison of platforms, browsers, and participant's devices (Preprint). *PsyArXiv*. https://doi.org/10.31234/osf.io/jfeca

Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23), 12750–12755. https://doi.org/10.1073/pnas.2002312117

Beckage, N., Todd, P. M., Penke, L., & Asendorpf, J. (2009). Testing sequential patterns in human mate choice using speed dating, In *Proceedings of the 2009 Cognitive Science Conference*.

Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, *138*(3), 389–414. https://doi.org/10.1037/a0026450

Cardinale, E. M., Pagliaccio, D., Swetlitz, C., Grassie, H., Abend, R., Costa, V. D., Averbeck, B. B., Brotman, M. A., Pine, D. S., Leibenluft, E., & Kircanski, K. (2021). Deliberative Choice Strategies in Youths: Relevance to Transdiagnostic Anxiety Symptoms. *Clinical Psychological Science*, 1–11. https://doi.org/10.1177/2167702621991805

Carpenter, T. P., Pogacar, R., Pullig, C., Kouril, M., Aguilar, S., LaBouff, J., Isenberg, N., & Chakroff, A. (2019). Survey-software implicit association tests: A methodological and empirical analysis. *Behavior Research Methods*, *51*(5), 2194–2208. https://doi.org/10.3758/s13428-019-01293-3

Casler, K., Bickel, L., & Hackett, E. (2013). Separate but equal? A comparison of participants and data gathered via Amazon's MTurk, social media, and face-to-face behavioral testing. *Computers in Human Behavior*, *29*(6), 2156–2160. https://doi.org/10.1016/j.chb.2013.05.009

Chambers, C., & Kording, K. (2018). While optimal models are useful, optimality claims are not that common. *PsyArXiv*. https://doi.org/10.31234/osf.io/6a35e

Chater, N., Zhu, J.-Q., Spicer, J., Sundh, J., León-Villagrá, P., & Sanborn, A. (2020). Probabilistic Biases Meet the Bayesian Brain. *Current Directions in Psychological Science*, *29*(5), 506–512. https://doi.org/10.1177/0963721420954801

Costa, V. D, & Averbeck, B. B. (2015). Frontal-Parietal and Limbic-Striatal Activity Underlies Information Sampling in the Best Choice Problem. *Cerebral Cortex*, *25*(4), 972–982. https://doi.org/10.1093/cercor/bht286

Craig, B. M., Hays, R. D., Pickard, A. S., Cella, D., Revicki, D. A., & Reeve, B. B. (2013). Comparison of US Panel Vendors for Online Surveys. *Journal of Medical Internet Research*, *15*(11), e260. https://doi.org/10.2196/jmir.2903

Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a Tool for Experimental Behavioral Research. *PloS one*, *8*(3). https://doi.org/10.1371/journal.pone.0057410

Dudey, T., & Todd, P. M. (2001). Making Good Decisions with Minimal Information: Simultaneous and Sequential Choice. *Journal of Bioeconomics*, *3*(2), 195–215. https://doi.org/10.1023/A:1020542800376

Eriksson, K., & Strimling, P. (2009). Partner search heuristics in the lab: Stability of matchings under various preference structures. *Adaptive Behavior*, *17*(6), 524–536. https://doi.org/10.1177/1059712309341220

Ford, J. B. (2017). Amazon's Mechanical Turk: A Comment. *Journal of Advertising*, *46*(1), 156–158. https://doi.org/10.1080/00913367.2016.1277380

Furl, N., & Averbeck, B. B. (2011). Parietal Cortex and Insula Relate to Evidence Seeking Relevant to Reward-Related Decisions. *Journal of Neuroscience*, *31*(48), 17572–17582. https://doi.org/10.1523/JNEUROSCI.4236-11.2011

Furl, N., Averbeck, B. B., & McKay, R. T. (2019). Looking for Mr(s) Right: Decision bias can prevent us from finding the most attractive face. *Cognitive psychology*, *111*, 1–14. https://doi.org/10.1016/j.cogpsych.2019.02.002

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, *61*(313), 35–73. https://doi.org/10.2307/2283044

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2017). Learning in the Repeated Secretary Problem. *arXiv preprint arXiv:1708.08831*, arxiv 1708.08831.

Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the Bayesians got their beliefs (and what those beliefs actually are): Comment on Bowers and Davis (2012). *Psychological Bulletin*, *138*(3), 415–422. https://doi.org/10.1037/a0026884

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, *114*(2), 211–244. https://doi.org/10.1037/0033-295X.114.2.211

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339–361. https://doi.org/10.1037/dec0000081

Guan, M., Lee, M. D., & Silva, A. (2014). Threshold Models of Human Decision Making on Optimal Stopping Problems in Different Environments. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*(36).

Guan, M., Lee, M. D., & Vandekerckhove, J. (2015). A Hierarchical Cognitive Threshold Model of Human Decision Making on Different Length Optimal Stopping Problems. *CogSci*.

Hauser, D., Paolacci, G., & Chandler, J. J. (2018). Common Concerns with MTurk as a Participant Pool: Evidence and Solutions. *PsyArXiv*.

Heit, E., & Erickson, S. (2011). In praise of secular Bayesianism. *Behavioral and Brain Sciences*, *34*(4), 169–231. https://doi.org/10.1017/S0140525X11000306

Hey, J. D. (1987). Still searching. *Journal of Economic Behavior & Organization*, *8*(1), 137–144. https://doi.org/10.1016/0167-2681(87)90026-6

Jones, M., & Love, B. C. (2011). Pinning down the theoretical commitments of Bayesian cognitive models. *Behavioral and Brain Sciences*, *34*(4), 215–231. https://doi.org/10.1017/S0140525X11001439

Kees, J., Berry, C., Burton, S., & Sheehan, K. (2017a). An Analysis of Data Quality: Professional Panels, Student Subject Pools, and Amazon's Mechanical Turk. *Journal of Advertising*, *46*(1), 141–155. https://doi.org/10.1080/00913367.2016.1269304

Kees, J., Berry, C., Burton, S., & Sheehan, K. (2017b). Reply to “Amazon’s Mechanical Turk: A Comment”. *Journal of Advertising*, *46*(1), 159–162. https://doi.org/10.1080/00913367.2017.1281781

Kemp, C., & Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychological Review*, *116*(1), 20–58.

Kochari, A. R. (2019). Conducting Web-Based Experiments for Numerical Cognition Research. *Journal of Cognition*, *2*(1), 1–21. https://doi.org/10.5334/joc.85

Lee, M. D. (2006). A Hierarchical Bayesian Model of Human Decision-Making on an Optimal Stopping Problem. *Cognitive Science*, *30*(3), 1–26. https://doi.org/10.1207/s15516709cog0000_69

Lee, M. D., O’Connor, T. A., & Welsh, M. B. (2005). Decision-Making on the Full Information Secretary Problem. *Proceedings of the Twenty-Sixth Conference of the Cognitive Science Society*, 819–824.

Marr, D. (1982). Vision. *San Francisco: Freeman*.

MATLAB. (2015). Version 2015b. The MathWorks Inc. *Natick, Massachusetts*.

Palan, S., & Schitter, C. (2018). Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, *17*, 22–27. https://doi.org/10.1016/j.jbef.2017.12.004

Paolacci, G., Chandler, J. J., & Ipeirotis, P. G. (2010). Running Experiments on Amazon Mechanical Turk. *Judgement and Decision making*, *5*(5), 411–419.

Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, *70*, 153–163. https://doi.org/10.1016/j.jesp.2017.01.006

Peterson, R. A. (2001). On the Use of College Students in Social Science Research: Insights from a Second-Order Meta-analysis. *Journal of Consumer Research*, *28*(3), 450–461. https://doi.org/10.1086/323732

Petzschner, F. H., Glasauer, S., & Stephan, K. E. (2015). A Bayesian perspective on magnitude estimation. *Trends in Cognitive Sciences*, *19*(5), 285–293. https://doi.org/10.1016/j.tics.2015.03.002

Prolific. (2014). Available at: https://www.prolific.co. *Oxford, UK*.

Samuel-Cahn, E. (1996). Optimal Stopping with Random Horizon with Application to the Full-Information Best-Choice Problem with Random Freeze. *Journal of the American Statistical Association*, *91*(433), 357–364. https://doi.org/10.1080/01621459.1996.10476696

Sanborn, A. N., & Chater, N. (2016). Bayesian Brains without Probabilities. *Trends in Cognitive Sciences*, *20*(12), 883–893. https://doi.org/10.1016/j.tics.2016.10.003

Sang, K., Todd, P. M., Goldstone, R. L., & Hills, T. T. (2020). Simple Threshold Rules Solve Explore/Exploit Trade-offs in a Resource Accumulation Search Task. *Cognitive Science*, *44*(2). https://doi.org/10.1111/cogs.12817

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

Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, *13*(4), 391–411. https://doi.org/10.1002/1099-0771(200010/12)13:4<391::AID-BDM359>3.0.CO;2-I

Seeley, T. D., & Buhrman, S. C. (2001). Nest-site selection in honey bees: How well do swarms implement the "best-of-N" decision rule? *Behavioral Ecology and Sociobiology*, *49*(5), 416–427. https://doi.org/10.1007/s002650000299

Sheehan, K. B. (2018). Crowdsourcing research: Data collection with Amazon's Mechanical Turk. *Communication Monographs*, *85*(1), 140–156. https://doi.org/10.1080/03637751.2017.1342043

Shi, L., Griffiths, T. L., Feldman, N. H., & Sanborn, A. N. (2010). Exemplar models as a mechanism for performing Bayesian inference. *Psychonomic Bulletin & Review*, *17*(4), 443–464. https://doi.org/10.3758/PBR.17.4.443

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. *Science*, *331*(6022), 1279–1285. https://doi.org/10.1126/science.1192788

Thomas, B., Coon, J., Westfall, H. A., & Lee, M. D. (2021). Model-Based Wisdom of the Crowd for Sequential Decision-Making Tasks. *Cognitive Science*, *45*(7). https://doi.org/10.1111/cogs.13011

Todd, P. M., & Gigerenzer, G. (2001). Putting naturalistic decision making into the adaptive toolbox. *Journal of Behavioral Decision Making*, *14*(5), 381–383. https://doi.org/10.1002/bdm.396

Todd, P. M., & Gigerenzer, G. (2003). Bounding rationality to the world. *Journal of Economic Psychology*, *24*(2), 143–165. https://doi.org/10.1016/S0167-4870(02)00200-3

Todd, P. M., & Miller, G. F. (1999). From pride and prejudice to persuasion: Satisficing in mate search, In *Simple heuristics that make us smart*. New York, NY, US, Oxford University Press.

Valone, T. J., Nordell, S. E., Giraldeau, L.-A., & Templeton, J. J. (1996). The empirical question of thresholds and mechanisms of mate choice. *Evolutionary Ecology*, *10*(4), 447–455. https://doi.org/10.1007/BF01237729

Chapter 3

Humans oversample in multiple decision-making domains

Abstract

Real-life sequential search problems such as finding a good parking spot require us to decide when to stop searching and accept the option currently available. Most studies on straightforward optimal stopping problems show that humans tend to terminate such searches too early compared to computational models of optimality. However, there are recent reports of search domains in which participants search too long instead, including our own findings using a facial attractiveness paradigm. Here, we explore whether such oversampling biases are an intrinsic feature of personal mate choice decisions, by investigating sampling decisions when participants choose attractive dates on a client's behalf (Study 1) and when they evaluate trustworthiness rather than attractiveness (Study 2). In both studies we replicated oversampling, again contradicting the consensus that undersampling biases are the norm. Furthermore, our results counter previous work that claims that oversampling biases arise because of the personal mate choice domain. We conclude that oversampling is not exclusively a feature of personal mate choice, but is a more general feature of decision making. Just how general remains to be determined, and more research is necessary to determine in which decision domains participants oversample or undersample and why.

3.1 Introduction

Decisions that require optimal stopping are remarkably common in everyday life. Individuals may encounter such sequentially presented search problems when choosing an apartment (Zwick et al., 2003), a parking space (Todd & Gigerenzer, 2012), or even a romantic partner (Furl et al., 2019). For decades, optimal stopping problems have been a popular field of study. To illustrate, a literature search in 2017 identified over 2000 papers published on the standard variant (the secretary problem, as defined by Ferguson, 1989) alone (Goldstein et al., 2017). The general consensus in the field of optimal stopping problems is that humans perform suboptimally compared to various models of optimality. More specifically, in a search for an option with the highest value when there is no explicit cost-to-sample, humans tend to terminate their searches too early. This observation will henceforth be referred to as *undersampling*. Undersampling has been reported on different optimal stopping tasks, including the secretary problem (Bearden et al., 2006; Seale & Rapoport, 1997), the beads task (Furl & Averbeck, 2011; Van der Leer et al., 2015), a numerical task (Guan et al., 2014), and full information tasks with different economic scenarios (Baumann et al., 2020; Cardinale et al., 2021; Costa & Averbeck, 2015).

Despite the evidence that humans show an undersampling bias on optimal stopping problems compared to models of optimality, there are recent reports of search domains in which participants search too long. A paper by Furl et al. (2019), for example, showed consistent evidence for an *oversampling* bias in a full information mate choice scenario called the facial attractiveness task. Like other full information problems described in the literature, the prior distribution was familiar to participants, there was no cost-to-sample, and outcomes were rewarded by their values (Abdelaziz & Krichen, 2006; Gilbert & Mosteller, 1966; Guan et al., 2014; Hill, 2009; Lee, 2006; Shu, 2008). However, the facial attractiveness task is different from the full information optimal stopping tasks that showed undersampling in at least three respects; 1) it investigates a different decision domain: mate choice, 2) it uses naturalistic image stimuli instead of abstract stimuli, and 3) participants themselves generate the prior distribution of option values, meaning that the value of a certain option is not set beforehand by the researcher and can differ between participants. More specifically, in the facial attractiveness task participants are instructed to make choices about whom to date. They do this by first rating the attractiveness of a pool of images, and then attempting to choose the most attractive face from a series of faces (each drawn from the initial pool) as their date.

Traditional mathematical treatments of optimal stopping problems have anecdotally described the problem in mate choice terms such as 'fiance' or 'dowry' problem (Gilbert & Mosteller, 1966). Yet the problem was not framed as mate choice in an empirical study in humans until Furl et al. (2019), who have created a contemporary paradigm that captures the essence of an online dating scenario (e.g., Tinder). This diverges from other mate choice paradigms that have attempted to incorporate 'real-life' elements such as rejection (Miller & Todd, 1998) and interaction (Eriksson & Strimling, 2009) in their tasks. Regardless of whether the facial attractiveness paradigm embodies every element that can occur in real-life dating, it is still possible that the mate choice domain was instrumental in changing participants' search strategies from undersampling to oversampling. Particularly, Furl et al. (2019) suggest that the design of the facial attractiveness task is sufficient to instigate mate choice predispositions. This mate choice bias, where individuals set high thresholds and continue searching for high-quality partners, has been observed across species including crickets (Ivy & Sakaluk, 2007), fiddler crabs (Backwell & Passmore, 1996), and sticklebacks (Milinski & Bakker, 1992). Participants in Furl et al. (2019) were observed to keep their thresholds too high throughout sequences, which consequently led to oversampling behaviour. However, if oversampling extends beyond the personal mate choice domain, then Furl et al. (2019)'s prediction that the facial attractiveness task instigates mate choice predispositions and their interpretation of participants' high thresholds in terms of mate choice is wrong.

Here, we report two novel studies of decision biases on full information optimal stopping tasks. The aim of both studies is to investigate whether oversampling biases are linked to the personal mate choice domain. Additionally, in our Supplementary Materials, we compare participants' sampling behaviour to four proposed computational models of sampling biases. In the first study (*matchmaker*), participants imagine that they are successful matchmakers. They are not making dating decisions for themselves, but for an imaginary 'client'. Because participants are not choosing their own mates, this design will allow us to explore whether oversampling extends to facial attractiveness in a broader sense. If oversampling is exclusively linked to the personal mate choice domain, participants in the matchmaker paradigm may abandon their irrational thresholds and make more objective and logical decisions when choosing attractive faces for a client. Consequently, participants may revert to optimal sampling or even the undersampling that is observed in economic domains. In contrast, if oversampling instead persists in this domain, then oversampling is not instigated specifically by personal mate choice, but arises from a more broadly operating mechanism. Finding this result would overturn Furl et al. (2019)'s explanation regarding participants' high thresholds, and instead would be in line with the literature proposing that the psychological mechanisms underlying facial attractiveness judgements are highly resistant adaptations that have evolved to identify information regarding an individual's health and mate quality in the general sense - not just for personal mate choice (Fink & Penton-Voak, 2002; Thornhill & Gangestad, 1999).

In the second study (*trustworthiness*), participants are asked to make decisions about the trustworthiness of the faces presented, rather than their attractiveness. If we discover in Study 1 that oversampling extends beyond the personal mate choice domain, then it is possible that oversampling extends even further into the domain of trustworthiness as well. However, while trustworthiness and attractiveness ratings have been found to be positively correlated (Bzdok et al., 2011; Todorov et al., 2008), they do not always uniformly affect decision-making (Ert et al., 2016; Jaeger et al., 2019). The literature suggests that attractiveness and trustworthiness are not interchangeable judgements and that the relationship between them is complex (Sofer et al., 2015; Stirrat & Perrett, 2010; Wilson & Eckel, 2006).

3.2 Materials and methods

3.2.1 Participants

Informed consent was obtained from all participants before the start of our studies, in accordance with the Declaration of Helsinki. Undergraduate students were recruited through the online Psychology Experiment Management System, as used by Royal Holloway, University of London. All participants received the same number of course credits for their participation, regardless of the time spent on the task. In line with the previously calculated effect size[1] for the facial attractiveness task (Furl et al., 2019), twenty undergraduate students aged 16-23 were enrolled in Study 1 (10 male participants and 10 female), and another twenty were enrolled in Study 2 (2 male participants and 18 female). Both studies were conducted in the lab at Royal

[1] Power analysis of Study 1 in Furl et al. (2019) suggested that fewer than 20 participants would be sufficient for 95% power.

Holloway, University of London, and ethical approval was granted by our institution.

3.2.2 Study 1: Matchmaker

The original design of the full information facial attractiveness task as described by Furl et al. (2019) was mostly maintained for the matchmaker paradigm, except for the instructions given to participants. Instead of making mate choice decisions for themselves, participants were instructed to imagine that they were successful matchmakers making mate choice decisions for a client (a friend, relative, or random person of the same sex). Participants were instructed that their attractiveness ratings and mate choice decisions should reflect the preferences of the imagined client, and not their own preferences.

Before starting the experiment, participants chose whether they wanted to rate male or female faces on their clients' behalf. In phase one, participants were asked to rate 426 unique faces on a scale from 1 (very unattractive, not someone your client would want to date) to 9 (very attractive, someone your client would want to date). The set of images used in this study was the same as the set used in Study 2 of Furl et al. (2019). The images were in colour and showed youthful individuals roughly between 18 and 30 years of age. Each face was rated three times, allowing for a final attractiveness rating for each face to be computed from the mean of these ratings. The reason for using each participant's own ratings to ultimately assess decision performance is that it eliminates the possibility of individual differences influencing the results, as both the participants and their corresponding computational models have the same prior of attractiveness ratings going into the sequence phase. In other words, phase one exposed both the participants and the model to the self-generated prior distribution of attractiveness values, from which the sequences in phase two are drawn (Furl et al., 2019).

Phase two of the study was the optimal stopping task. Participants were told that the faces they would encounter were randomly sampled from the pool of faces that were rated in phase one. In other words, for each individual participant, phase two faces were generated from a random subset of phase one faces, allowing for variable distributions of attractiveness values between sequences. Participants were not given any explicit information regarding the attractiveness distribution that generated the sequences, except that the sequence options would be drawn from the faces rated during phase one. Participants were shown 28 sequences of 8 faces and reminded to think like their client - which face would their client think is the most attractive and would they like to go on a date with? The conditions of the task were that they could not return to a previously rejected option, and that it was impossible to know for certain the value of the upcoming option(s) in any given sequence. The number of options remaining was shown at the top of the screen, and the rejected options were shown at the bottom of the screen. When a participant made a choice, they had to advance through a series of grey squares that replaced the remaining faces. This ensured that participants could not finish the study early by choosing an early option. Adding grey squares does not alter participants' sampling behaviour: Furl et al., 2019 found the same results on the facial attractiveness paradigm with the implementation of grey squares (Studies 2 and 3) and without (Study 1). The entire study was self-paced. When a participant reached the end of a sequence without having made a choice, the last option in the sequence became their client's date by default. The two key dependent variables of interest are the number of samples before choice (i.e., the position of the chosen image in the sequence), and the rank of the chosen option. Ranks were computed for every sequence according to each participant's own attractiveness ratings, with rank 1 corresponding to the lowest rated face and rank 8 corresponding to the highest rated face in that sequence. Both key variables are a mean value over the sequences for each participant.

3.2.3 Study 2: Trustworthiness

The trustworthiness paradigm was very similar to both the facial attractiveness task described by Furl et al. (2019) and the matchmaker paradigm described in Section 3.2.2. Here, we will only highlight the differences between the trustworthiness paradigm and the matchmaker paradigm.

In phase one of the trustworthiness paradigm, participants rated faces on their perceived trustworthiness. Participants rated on a scale from one to nine, with one indicating very untrustworthy and nine indicating very trustworthy. The same set of faces was used as previously in the matchmaker paradigm.

Next, participants commenced with the sequence part of the study in which they were presented with 28 sequences of eight faces. The goal was to pick as trustworthy a face as possible in each sequence under the conditions that it was impossible to go back to a previously shown face, and it was impossible to know the value of the face(s) yet to come.

3.2.4 Data analysis

We examined participants' sampling biases by comparing their sampling behaviour to that of a Bayesian ideal observer model (Costa & Averbeck, 2015) where performance is Bayesian optimal, using MATLAB version 2015b (MATLAB, 2015). The Bayesian ideal observer model extends the framework for full information problems described by Gilbert and Mosteller (1966), which computes the expected values of future options under a standard normal distribution via a backwards induction technique. The original model by Gilbert and Mosteller (1966) knows for all sequence options the mean and variance of the distribution of option values from which each option value in a sequence is drawn (i.e., the generating distribution), and that the reward values of the sequence ranks are equal to the option values. The Bayesian nature of the ideal observer model further allows the mean and variance of the (prior) distribution to be updated and learned as sequences progress, based on the newly sampled option values in the sequences. Moreover, it adds a cost-to-sample parameter and allows outcome ranks to be weighted by a reward function. We did not use monetary incentivisation in our task and instead, assumed that participants followed our instructions by maximising the ranks of their choices. We therefore used a reward function similar to that traditionally used in Gilbert and Mosteller (1966)'s version of the model, in which the ideal observer agent is rewarded proportional to the value of the chosen options. Because there was no extrinsic sample cost in our paradigm, we fixed the cost-to-sample parameter to zero, while other model values were fixed to the values assigned by Furl et al. (2019). The ideal observer model received as input for each participant the same option values of the sequences as presented to the participant in phase two, with the participant's individual rating of each option (averaged over the three ratings provided) as the sequence values. To ensure normality, ratings were log transformed for each participant before being entered into the model.

Statistical tests were performed using RStudio (RStudioTeam, 2020). For all analyses, a p value of $< .05$ was considered significant. Additionally, to provide a compelling answer to our research question, we show the Bayes factor for statistical tests of the mean number of samples and mean rank as well. A particular advantage of reporting Bayes factor is that it allows us to make a statement about the strength of the evidence in favour of either the null hypothesis or the alternative hypothesis (Jarosz & Wiley, 2014). Bayesian *t*-tests comparing participants and the ideal observer model were calculated using the BayesFactor package (Morey & Rouder, 2018), within the R environment (RStudioTeam, 2020).

3.3 Results

3.3.1 Oversampling bias in Study 1 and Study 2

Initial evaluation of the data revealed that one participant in Study 2 had not sufficiently engaged with the task, as they had failed to make a choice on any of the sequences and always sampled until the final option, which they were then forced to choose. This participant was therefore excluded from further analysis.

First, we tested for effects of faces' sex, as chosen by participants. For Study 1, a one-way ANOVA showed no effect of chosen sex on the number of samples ($F(1, 18) = 0.050$, $p = .826$), nor on the ranks of the chosen faces ($F(1, 18) = 0.546$, $p = .470$). Fourteen participants rated male faces in Study 1, and six participants rated female faces. For Study 2, a one-way ANOVA showed no effect of chosen sex on the number of samples ($F(1, 17) = 1.697$, $p = .210$), nor on the ranks of the chosen faces ($F(1, 17) = 0.735$, $p = .403$). Five participants rated male faces in Study 2, and fourteen participants rated female faces.

Participants in both studies showed an oversampling bias ($M_1 = 5.31$, $SD_1 = 0.84$, $M_2 = 5.60$, $SD_2 = 0.66$) compared to the Bayesian ideal observer model ($M_1 = 3.73$, $SD_1 = 0.54$, $M_2 = 3.53$, $SD_2 = 0.54$). Two-tailed *t*-tests pairing participants with their corresponding ideal observer models, indicated that participants sampled more and ended up with lower-ranked options than the optimal model (Table 3.1). For Study 2, the non-parametric Wilcoxon Signed-Rank test was used to compute the differences in mean rank as this variable was not normally distributed for participants. We note that the matchmaker and trustworthiness paradigms yielded similar results. Bayes factors for the mean number of samples and mean rank are shown in Table 3.2.

TABLE 3.1: Pairwise two-tailed t-tests for differences between participants and the ideal observer model for mean number of samples and mean rank of the chosen option. For Study 2, a Wilcoxon Signed-Rank test was used to compute the differences in mean rank. Also reported are Cohen's d effect sizes.

		Ideal Observer	
		Mean number of samples	Mean rank
Participants	Study 1	d = -1.70 $t(19)$ = -7.60 p <.001	d = 1.57 $t(19)$ = 7.03 p <.001
	Study 2	d = -2.43 $t(17)$ = -10.32 p <.001	d = 1.32 Z = 166 p <.001

TABLE 3.2: Bayes factor (BF_{10}) describing the relative likelihood of a difference between participants and the ideal observer model for mean number of samples and mean rank of the chosen option. $BF_{10} > 100$ can be interpreted as extreme or decisive evidence for the alternative hypothesis (Jarosz & Wiley, 2014; Wagenmakers et al., 2018), i.e., there is a difference between participants and the ideal observer model.

		Ideal Observer	
		Mean number of samples	Mean rank
Participants	Study 1	$BF_{10} = 4.589e^4 \pm 0\%$	$BF_{10} = 1.678e^4 \pm 0\%$
	Study 2	$BF_{10} = 1.259e^6 \pm 0\%$	$BF_{10} = 7.977e^2 \pm 0\%$

3.3.2 Comparison of rating values in Study 1 and Study 2

Having confirmed that participants showed an equivalent oversampling bias in Study 1 and Study 2, we wanted to investigate whether participants judged attractiveness and trustworthiness in a similar way. However, even though methodologically there is little difference between the matchmaker paradigm and the trustworthiness paradigm other than the instructions given to participants, a remarkable disparity was observed in the subjective rating values. When a density plot was used to visualise the distributions of participants' ratings (i.e., the generating distribution of option values), we noticed a clear difference between attractiveness and trustworthiness ratings (Figure 3.1). As stated in the methods section, participants rated the same set of faces in both studies. Nevertheless, the distribution of attractiveness ratings appears to have a lower mean and a more positive skew, while trustworthiness ratings appear more normally distributed with a higher mean. The apparent difference between the matchmaker and trustworthiness distributions was confirmed with a Kolmogorov-Smirnov test: $D(426) = 0.75$, $p < .001$ (Bonferroni corrected for the three distributions, as shown below).

To explore whether participants' ratings in our Study 1 and Study 2 resembled participants' attractiveness ratings in Furl et al. (2019)'s facial attractiveness task, we also plotted in Figure 3.1 the distribution of rating values for Study 2 of Furl et al. (2019) (a more detailed description of this study can be found in our Supplementary Materials). We observe that the personal mate choice ratings as recorded by Furl et al. (2019) best resembled our matchmaker task, particularly in terms of the mean of the distribution, but differences in the shape of the distribution can nonetheless be perceived (e.g., skewness). We confirmed this difference between the matchmaker and attractiveness distributions with a Kolmogorov-Smirnov test: $D(426) = 0.16$, $p < .001$ (Bonferroni corrected for the three distributions). This finding, in combination with our observations of Figure 3.1, is compelling evidence that the three studies plotted here capture three different sets of judgements of the same set of faces, but nevertheless oversampling biases are observed in each study.

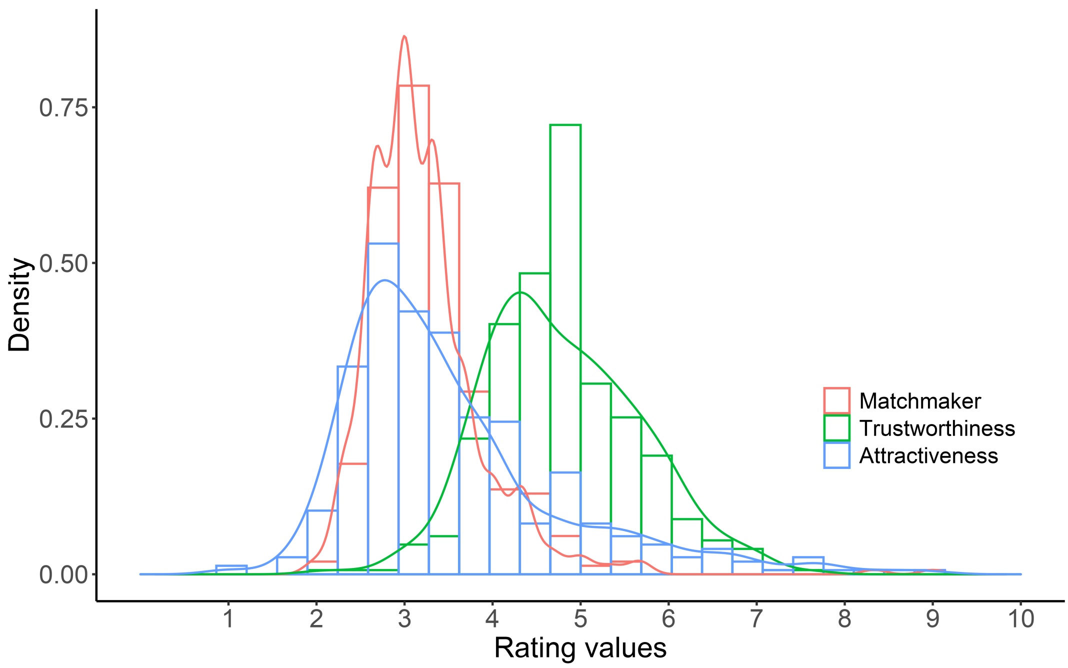

FIGURE 3.1: Density plot with overlaying histogram showing the distributions of participants' matchmaker attractiveness ratings (Matchmaker) and trustworthiness ratings (Trustworthiness), as well as facial attractiveness ratings (Attractiveness) as collected in Furl et al. (2019), Study 2.

3.4 Discussion

Previous literature on optimal stopping problems has repeatedly reported an undersampling bias in human participants compared to mathematical definitions of optimality (Baumann et al., 2020; Bearden et al., 2006; Cardinale et al., 2021; Costa & Averbeck, 2015; Furl & Averbeck, 2011; Guan et al., 2014; Seale & Rapoport, 1997). The aim of the studies reported in this paper was to further investigate the contradictory oversampling bias found by Furl et al. (2019) in a mate choice scenario by looking at variations of the full information facial attractiveness task. The first study utilised a matchmaker paradigm. Participants who were tasked to find the perfect date for a hypothetical client sampled more and ended up with lower-ranked options compared to a Bayesian ideal observer model of optimality. This indicated that oversampling is not a function of personal mate choice specifically, but perhaps is associated with facial attractiveness more generally. Our second study – which utilised a trustworthiness paradigm – suggests an even wider scope: participants also oversampled when making decisions about the trustworthiness of faces. That is, compared to a Bayesian ideal observer model, participants sampled more and ended up with lower-ranked options. Based on the results of our Bayes Factor analysis we can state that there is better evidence for our alternative hypothesis (participants' sampling behaviour differs from that of a Bayesian ideal observer model) than the null hypothesis (of no difference in sampling behaviour between participants and the Bayesian ideal observer model) (Jarosz & Wiley, 2014; Wagenmakers et al., 2018).

An explanation for the oversampling found in Study 1 could be that the matchmaker paradigm is not a sufficiently different task from the facial attractiveness task in Furl et al. (2019). Indeed, to the extent that participants used their own preferences as a proxy for their fictional client, it is conceivable that personal mate choice might still introduce some oversampling. Although not all participants chose to rate faces of the opposite sex, participants' sexual preferences were not collected so we cannot draw any conclusions regarding whether participants made decisions congruent with their own personal mate choice. However, this explanation remains inadequate because we also found oversampling in Study 2, and trustworthiness decisions are even less likely to instigate mate choice predispositions than matchmaker decisions. Some literature suggests that characteristics including warmth and potential trustworthiness are criteria for mate selection (Fletcher et al., 1999; Fletcher et al., 2004; Valentine et al., 2020). However, our trustworthiness paradigm did not mention the words 'date' or 'dating', and as such made no suggestion towards any form of mate selection, unlike the facial attractiveness task. Therefore, it seems unlikely that participants implicitly connected rating faces on trustworthiness to finding a potential romantic partner. Furthermore, a clear disparity could be observed between the distribution of attractiveness ratings in Study 1 and the distribution of trustworthiness ratings in Study 2 (Figure 3.1), which implies that participants used at least some different visual information to judge trustworthiness than to judge what another person might find attractive. This argument is supported by literature on face typicality, which has shown that the typical face is perceived as the most trustworthy, but not the most attractive (Sofer et al., 2015). These findings are in line with our participants' rating distributions (Figure 3.1). From this, we infer that the trustworthiness paradigm did not instigate mate choice predisposition, and consequently that mate choice cannot explain oversampling biases.

Our findings contradict one of the principal conclusions of Furl et al. (2019), by indicating that the mate choice domain is not the regulating factor for the conflicting oversampling bias observed on the facial attractiveness task. Instead, our results suggest that the tendency to sample more than is optimal spans different decision-making domains, including trustworthiness. This also sheds new light on the high threshold theory proposed by Furl et al. (2019) as an explanation for oversampling. The high threshold theory stems from the behavioural ecology literature on mate choice, where individuals set high thresholds and continue searching for high-quality partners (Backwell & Passmore, 1996; Ivy & Sakaluk, 2007; Milinski & Bakker, 1992). As our findings indicate that oversampling extends beyond the mate choice domain, the oversampling bias observed on the facial attractiveness task must not have been for this reason, and Furl et al. (2019)'s conclusions should now be reevaluated (for a computational model comparison, see Supplementary Materials).

In contrast to the tasks that have yielded undersampling, both the original facial attractiveness task and the present adaptations of it used naturalistic image stimuli instead of abstract stimuli (i.e., numbers, text). Considering that so many real-world decisions depend on searching through natural images, it is important for future research to investigate whether the type of stimuli (images or otherwise) used in an optimal stopping task is predictive of the sampling bias shown by participants. Ultimately, more decision-making domains, including image-based ones, remain to be tested before a conclusive answer regarding sampling biases on optimal stopping problems can be obtained.

To conclude, our findings are an important corrective to the consensus view that humans undersample on (full information) optimal stopping tasks. They are also an important corrective to Furl et al. (2019)'s conclusion that personal mate choice has a special ability to lead to oversampling biases. Instead, we show that oversampling can arise across different decision-making domains, thus revealing that oversampling is a much more general phenomenon than previously believed. Just how general remains to be determined, and more research is necessary to determine in which decision domains participants oversample or undersample and why.

References

Abdelaziz, F. B., & Krichen, S. (2006). Optimal stopping problems by two or more decision makers: A survey. *Computational Management Science*, *4*(2), 89–111. https://doi.org/10.1007/s10287-006-0029-5

Backwell, P. R. Y., & Passmore, N. I. (1996). Time constraints and multiple choice criteria in the sampling behaviour and mate choice of the fiddler crab, Uca annulipes. *Behavioral Ecology and Sociobiology*, *38*(6), 407–416. https://doi.org/10.1007/s002650050258

Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23), 12750–12755. https://doi.org/10.1073/pnas.2002312117

Bearden, J. N., Rapoport, A., & Murphy, R. (2006). Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Science*, *52*(9), 1437–1449. https://doi.org/10.1287/mnsc.1060.0535

Bzdok, D., Langner, R., Caspers, S., Kurth, F., Habel, U., Zilles, K., Laird, A., & Eickhoff, S. B. (2011). ALE meta-analysis on facial judgments of trustworthiness and attractiveness. *Brain Structure and Function*, *215*(3), 209–223. https://doi.org/10.1007/s00429-010-0287-4

Cardinale, E. M., Pagliaccio, D., Swetlitz, C., Grassie, H., Abend, R., Costa, V. D., Averbeck, B. B., Brotman, M. A., Pine, D. S., Leibenluft, E., & Kircanski, K. (2021). Deliberative Choice Strategies in Youths: Relevance to Transdiagnostic Anxiety Symptoms. *Clinical Psychological Science*, 1–11. https://doi.org/10.1177/2167702621991805

Costa, V. D, & Averbeck, B. B. (2015). Frontal-Parietal and Limbic-Striatal Activity Underlies Information Sampling in the Best Choice Problem. *Cerebral Cortex*, *25*(4), 972–982. https://doi.org/10.1093/cercor/bht286

Eriksson, K., & Strimling, P. (2009). Partner search heuristics in the lab: Stability of matchings under various preference structures. *Adaptive Behavior*, *17*(6), 524–536. https://doi.org/10.1177/1059712309341220

Ert, E., Fleischer, A., & Magen, N. (2016). Trust and reputation in the sharing economy: The role of personal photos in Airbnb. *Tourism Management*, *55*, 62–73. https://doi.org/10.1016/j.tourman.2016.01.013

Ferguson, T. S. (1989). Who Solved the Secretary Problem? *Statistical Science*, *4*(3), 282–289. https://doi.org/10.1214/ss/1177012493

Fink, B., & Penton-Voak, I. (2002). Evolutionary Psychology of Facial Attractiveness. *Current Directions in Psychological Science*, *11*(5), 154–158. https://doi.org/10.1111/1467-8721.00190

Fletcher, G. J. O., Simpson, J. A., Thomas, G., & Giles, L. (1999). Ideals in intimate relationships. *Journal of Personality and Social Psychology*, *76*(1), 72–89. https://doi.org/10.1037/0022-3514.76.1.72

Fletcher, G. J. O., Tither, J. M., O'Loughlin, C., Friesen, M., & Overall, N. (2004). Warm and Homely or Cold and Beautiful? Sex Differences in Trading Off Traits in Mate Selection. *Personality and Social Psychology Bulletin*, *30*(6), 659–672. https://doi.org/10.1177/0146167203262847

Furl, N., & Averbeck, B. B. (2011). Parietal Cortex and Insula Relate to Evidence Seeking Relevant to Reward-Related Decisions. *Journal of Neuroscience*, *31*(48), 17572–17582. https://doi.org/10.1523/JNEUROSCI.4236-11.2011

Furl, N., Averbeck, B. B., & McKay, R. T. (2019). Looking for Mr(s) Right: Decision bias can prevent us from finding the most attractive face. *Cognitive psychology*, *111*, 1–14. https://doi.org/10.1016/j.cogpsych.2019.02.002

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, *61*(313), 35–73. https://doi.org/10.2307/2283044

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2017). Learning in the Repeated Secretary Problem. *arXiv preprint arXiv:1708.08831*, arxiv 1708.08831.

Guan, M., Lee, M. D., & Silva, A. (2014). Threshold Models of Human Decision Making on Optimal Stopping Problems in Different Environments. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*(36).

Hill, T. P. (2009). Knowing When to Stop: How to gamble if you must—the mathematics of optimal stopping. *American Scientist*, *97*(2), 126–133.

Ivy, T. M., & Sakaluk, S. K. (2007). Sequential mate choice in decorated crickets: Females use a fixed internal threshold in pre- and postcopulatory choice. *Animal Behaviour*, *74*(4), 1065–1072. https://doi.org/10.1016/j.anbehav.2007.01.017

Jaeger, B., Sleegers, W. W. A., Evans, A. M., Stel, M., & van Beest, I. (2019). The effects of facial attractiveness and trustworthiness in online peer-to-peer markets. *Journal of Economic Psychology*, *75*, 102125. https://doi.org/10.1016/j.joep.2018.11.004

Jarosz, A. F., & Wiley, J. (2014). What Are the Odds? A Practical Guide to Computing and Reporting Bayes Factors. *The Journal of Problem Solving*, *7*(1). https://doi.org/10.7771/1932-6246.1167

Lee, M. D. (2006). A Hierarchical Bayesian Model of Human Decision-Making on an Optimal Stopping Problem. *Cognitive Science*, *30*(3), 1–26. https://doi.org/10.1207/s15516709cog0000_69

MATLAB. (2015). Version 2015b. The MathWorks Inc. *Natick, Massachusetts*.

Milinski, M., & Bakker, T. (1992). Costs influences sequential mate choice in sticklebacks, gasterosteus aculeatus. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *250*(1329), 229–233. https://doi.org/10.1098/rspb.1992.0153

Miller, G. F., & Todd, P. M. (1998). Mate choice turns cognitive. *Trends in Cognitive Sciences*, *2*(5), 190–198. https://doi.org/10.1016/S1364-6613(98)01169-3

Morey, R. D., & Rouder, J. N. (2018). BayesFactor: Computation of Bayes Factors for Common Designs.

RStudioTeam. (2020). RStudio: Integrated Development Environment for R. *Boston, Massachusetts*.

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

Shu, S. B. (2008). Future-biased search: The quest for the ideal. *Journal of Behavioral Decision Making*, *21*(4), 352–377. https://doi.org/10.1002/bdm.593

Sofer, C., Dotsch, R., Wigboldus, D. H. J., & Todorov, A. (2015). What Is Typical Is Good: The Influence of Face Typicality on Perceived Trustworthiness. *Psychological Science*, *26*(1), 39–47. https://doi.org/10.1177/0956797614554955

Stirrat, M., & Perrett, D. (2010). Valid Facial Cues to Cooperation and Trust: Male Facial Width and Trustworthiness. *Psychological Science*, *21*(3), 349–354. https://doi.org/10.1177/0956797610362647

Thornhill, R., & Gangestad, S. W. (1999). Facial attractiveness. *Trends in Cognitive Sciences*, *3*(12), 452–460. https://doi.org/10.1016/S1364-6613(99)01403-5

Todd, P. M., & Gigerenzer, G. (2012). *Ecological Rationality Intelligence in the World*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195315448.001.0001

Todorov, A., Said, C. P., Engell, A. D., & Oosterhof, N. N. (2008). Understanding evaluation of faces on social dimensions. *Trends in Cognitive Sciences*, *12*(12), 455–460. https://doi.org/10.1016/j.tics.2008.10.001

Valentine, K. A., Li, N. P., Meltzer, A. L., & Tsai, M.-H. (2020). Mate Preferences for Warmth-Trustworthiness Predict Romantic Attraction in the Early Stages of Mate Selection and Satisfaction in Ongoing Relationships. *Personality and Social Psychology Bulletin*, *46*(2), 298–311. https://doi.org/10.1177/0146167219855048

Van der Leer, L., Hartig, B., Goldmanis, M., & McKay, R. (2015). Delusion proneness and 'jumping to conclusions': Relative and absolute effects. *Psychological Medicine*, *45*(6), 1253–1262. https://doi.org/10.1017/S0033291714002359

Wagenmakers, E.-J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Selker, R., Gronau, Q. F., Dropmann, D., Boutin, B., Meerhoff, F., Knight, P., Raj, A., van Kesteren, E.-J., van Doorn, J., Šmíra, M., Epskamp, S., Etz, A., Matzke, D., . . . Morey, R. D. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic Bulletin & Review*, *25*(1), 58–76. https://doi.org/10.3758/s13423-017-1323-7

Wilson, R., & Eckel, C. (2006). Judging a book by its cover: Beauty and expectations in the trust game. *Political Research Quarterly*, *59*(2), 189–202. https://doi.org/10.1177/106591290605900202

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer Sequential Search: Not Enough or Too Much? *Marketing Science*, *22*(4), 503–519. https://doi.org/10.1287/mksc.22.4.503.24909

Supplementary materials

S1 Background

Furl et al. (2019) found that the oversampling bias on the facial attractiveness task was best described by the *biased values model* (BV): a model of human behaviour which gives less weighting to low-quality options to explain the lack of substantial threshold decline. The biased values model was motivated by high threshold theories from behavioural ecology, as outlined in the Introduction of our main manuscript (Backwell & Passmore, 1996; Ivy & Sakaluk, 2007; Milinski & Bakker, 1992). Furl et al. (2019)'s conclusion that the mate choice domain should selectively lead to oversampling is, in part, based on the success of the *biased values model* in their model comparison. However, our studies show that oversampling is not limited to the personal mate choice domain, and as such, it is worth revisiting the model comparison for our new datasets to see whether there is equally strong evidence for the biased values model.

Other bias models considered by Furl et al. (2019) to explain participants' oversampling behaviour were the *sample reward model* (SR) and the *attractive prior model* (AP). The sample reward model implements an intrinsic reward for sampling, while the attractive prior model implements a maladaptively high mean value for the prior distribution. Although the biased values model correlated best with participants' pattern of choices, both the sample reward model and the attractive prior model resembled participants' data in the sense that they oversampled, compared to the Bayesian ideal observer model (Furl et al., 2019). As such, the sample reward model and attractive prior model were at least better correlated with participants' sampling behaviour than the optimal model, and could explain participants' oversampling biases in Furl et al. (2019) to some degree.

Another contending model of human decision-making proposed in the optimal stopping literature is the *cutoff model* (CO) (Dudey & Todd, 2001; Sang et al., 2020; Seale & Rapoport, 1997). This model, first suggested by Kahan et al. (1967), samples a certain number of options (the cutoff proportion), determines what the best option was in this sample, and sets this as the cutoff threshold. It then continues sampling until an option is encountered that exceeds the cutoff threshold, or until the end of the sequence is reached in which case the last option is chosen. The cutoff model does not use a reward function, but instead assumes that the agent is only interested in choosing the top-ranked value and there is no cost-to-sample. However, despite this assumption, the model is also accurate in practice at maximising the average rank of the chosen option, and therefore, is at least somewhat robust to relaxations of that assumption (Todd & Gigerenzer, 2003; Todd & Miller, 1999). The cutoff model was found to best describe participants' sampling behaviour on a classic optimal stopping task: the standard 'Secretary Problem' (Seale & Rapoport, 1997). Dudey and Todd (2001) replicated these findings in a simulation study and showed that the cutoff model outperformed other models. Furthermore, the cutoff model was found to generalise to some extent to other decision-making scenarios. On an explore/exploit trade-off task where both situations (exploration and exploitation) were rewarding, Sang et al. (2020) found that the cutoff model was the second best fitting model (after the 'linear decreasing threshold model'), describing decision-making in 20% of participants. However, other authors have found no support for the cutoff model, rejecting its usefulness for describing human behaviour and declaring it suitable only for the simplest of optimal stopping paradigms (Baumann et al., 2020; Lee et al., 2005).

In these supplementary analyses we further investigate the three models of human decision-making biases proposed by Furl and colleagues (2019), as well as the cutoff model. All models are considered here to be viable candidates for explaining human sampling behaviour. Whilst previous research on the facial attractiveness task has favoured the biased values model (Furl et al., 2019), this could be due to the biased values model embodying inherent mate choice predispositions that may not extend to different decision making domains. By including the cutoff model in our analysis, we aim to contribute to the debate regarding the cutoff model as a model of human decision making on optimal stopping tasks. We include the *Bayesian ideal observer model* (IO) as well in our comparison, although based on our main results we do not expect this model to be a good fit for describing human sampling behaviour.

S2 Data analysis

For each study, we fitted hypothetical models to participants' sampling behaviour, relying on procedures described in Furl et al. (2019), Costa and Averbeck (2015), Cardinale et al. (2021) and Furl and Averbeck (2011). The key computation for each model was utility (equation 3.1), which depends on conditional probabilities and reward values (equation 3.2), which are based on how each option ranks relative to the other options in its sequence. Models use backwards induction to derive utilities that could result from further sampling (equation 3.3).

$$u_t\left(s_t\right) = \max_{a \in A_{s_t}} \left\{ r_t\left(s_t, a\right) + \int_s p_t\left(j \mid s_t, a\right) u_{t+1}(j) dj \right\} \tag{3.1}$$

The utility u of the state s at sample t is the value of the best action a, which depends on reward value r, the cost-to-sample C_s, and the probabilities of outcomes j of subsequent states, weighted by their utilities.

$$\begin{aligned} r_t\left(s_t, a = accept\right) &= \sum_{i=1}^{N} p\left(rank = i\right) * R\left(i + (h-1)\right) \\ r_t\left(s_t, a = decline\right) &= C_s \end{aligned} \tag{3.2}$$

Because we are interested in how participants make decisions on the basis of intrinsic social reward values (and not extrinsic economic ones), we did not provide extrinsic monetary incentivisation for choosing highly rated options. Instead, we instructed participants to choose the option with the highest-ranked value possible. As with most non-incentivised tasks used in behavioural studies, participants were clearly able to achieve quite accurate performance based on their own intrinsic reward functions, without the need for additional extrinsic incentives. We formalised participants' intrinsic reward functions by simply assuming they followed instructions and attempted to maximise the rank of the option chosen. We therefore assigned to outcome reward r the relative rank h for each option i in a sequence, as well as the cost-to-sample C_s. When considering final sequence position N, the model computes final utilities as:

$$u_N\left(s_N\right) = r\left(s_N\right) \text{ for all } s_N \in N \tag{3.3}$$

and working backwards from N, we use equation 3.1 to compute utilities at every sequence position t.

To calculate the value of either choosing or declining an option in a sequence we compute the action value Q as:

$$\begin{aligned} Q_t\left(s_t, a = take\right) &= r_t\left(s_t, a\right) \\ Q_t\left(s_t, a = decline\right) &= \int_s p_t\left(j \mid s_t, a\right) u_{t+1}(j) dj \end{aligned} \tag{3.4}$$

We used a softmax transformation (equation 3.5) to normalise the computational models' action values to probabilities, and we fit the softmax sensitivity parameter β as an additional free parameter in all models, with smaller β values producing more stochasticity (Baumann et al., 2020). The individual models described in the next section were all fit in this way. The softmax equation with the action values for taking an option Q_{take}, and the action values for declining an option $Q_{decline}$ is defined as:

$$Softmax(Q_t) = \left[\frac{e^{Q_{take}}}{\sum e^{Q_{decline}}}\right] \tag{3.5}$$

To create the biased computational models, the input and/or model values of the ideal observer model were altered, in an effort to have the ideal observer model reproduce participants' biases. We will describe these computational models in more detail in the next paragraph. It is expected that the computational models that include free parameters to explain bias provide a better match to participants' behaviour than the ideal observer model. To compute these free parameters, we used an optimisation search algorithm.

Our first computational model, the biased values model, transforms attractiveness values using a logistic utility function that effectively limits the influence of lower-rated faces on sampling behaviour, thus raising the choice threshold of the model. The biased values model resembles participants who make decisions using Bayesian optimal computations, except that their computations operate on 'misperceived' option values. Hence, mathematically, the biased values model is the same as the ideal observer model described above. The free parameters for the biased values model are the slope and midpoint of the logistic utility function, as well as the β value mentioned above. The mean and the standard deviation of the estimated parameter values for all five of our computational models, across our three datasets, are shown in Table S1.

Our second computational model is the sample reward model. As explained in our main manuscript, the cost-to-sample parameter C_s was set to zero for the ideal observer model. In contrast, we optimised the cost-to-sample parameter in the sample reward model, thus introducing an intrinsic aversion

TABLE S1: Mean and standard deviation of the estimated parameter values for the biased values model (BV), sample reward model (SR), attractive prior model (AP), cutoff model (CO), and ideal observer model (IO), across our three datasets: matchmaker, trustworthiness and attractiveness.

Model	Parameters	Matchmaker	Trustworthiness	Attractiveness
BV	Slope	M = 1.319, SD = 0.785	M = 0.884, SD = 0.618	M = 1.001, SD = 0.108
	Midpoint	M = 4.701, SD = 3.637	M = 14.647, SD = 10.142	M = 10.592, SD = 5.693
	β	M = 2.828, SD = 2.626	M = 1.577, SD = 0.401	M = 1.568, SD = 0.311
SR	Cost-to-sample	M = -0.021, SD = 0.084	M = -0.072, SD = 0.289	M = -0.026, SD = 0.068
	β	M = 5.047, SD = 5.247	M = 11.597, SD = 15.353	M = 7.643, SD = 4.098
AP	Constant	M = 0.526, SD = 0.234	M = 0.428, SD = 0.194	M = 0.373, SD = 0.297
	β	M = 5.324, SD = 6.884	M = 7.121, SD = 4.023	M = 7.825, SD = 3.961
CO	Cutoff	M = 3.571, SD = 0.866	M = 3.341, SD = 0.819	M = 3.649, SD = 0.705
	β	M = 0.818, SD = 0.545	M = 1.179, SD = 0.537	M = 1.234, SD = 0.263
IO	β	M = 0.000, SD = 0.000	M = 0.000, SD = 0.000	M = 0.000, SD = 0.000

(positive cost-to-sample value) or attraction (negative cost-to-sample value) to sampling. Also added as a free parameter to the sample reward model is the β value. Mathematically, the sample reward model is the same as the ideal observer model described above.

Thirdly, for the attractive prior model, we optimised a constant which was added to the prior mean, with the expectation that it would turn out positive because this would result in oversampling. Namely, by adding a positive constant to the prior mean, a maladaptively high mean value is created which leads to a holding-out effect. The attractive prior model thus resembles optimism biases in the distribution of prior beliefs. Also added as a free parameter to the attractive prior model is the β value. Mathematically, the attractive prior model is the same as the ideal observer model described above.

For the ideal observer model, from which the biased values, sample reward and attractive prior models are derived, the only free parameter is the β value.

Finally, recall that the cutoff model samples a certain number of options during a 'learning period' after which the next option is chosen that exceeds the cutoff threshold (i.e., the value of the best option seen during the learning period). The length of the learning period is the cutoff parameter. Also added as a free parameter to the cutoff model is the β value. Mathematically, the cutoff model can be described by two integers (r, s), where r is the number of options seen so far, and s is the rank of the rth, last presented option (Seale & Rapoport, 1997). After the next option has been presented, the new state of the cutoff model will be $(r + 1, s')$, where s' is equally likely to be any one of the integers 1, 2, ..., r + 1. The probability that an option is the best of all the n options is r/n. The equation for the maximum probability of choosing the best option is

$$a_r = 1/r + 1/(r+1) + ... + 1/(n-1) \tag{3.6}$$

The cutoff model then rejects the first r - 1 options after which it chooses the next top-ranked option (Seale & Rapoport, 1997).

We optimised the free parameters of each model to minimise the negative log likelihood (i.e., best model fit). As starting values for the parameters we used $1/e$ = 0.368 (the optimal cutoff proportion, Gilbert & Mosteller, 1966; Seale & Rapoport, 1997) for the cutoff model and, for the Bayesian models, we used starting values taken from the previous analysis of facial attractiveness described in Furl et al. (2019). Next, the Akaike Information Criterion (AIC) was used to estimate how well each model fitted the data, while correcting for the number of free parameters. Note that the AIC score is minimised, i.e., a lower AIC score means a better fit of the model to the data. The 'best' model (a possible theory of human behaviour) should not only yield the best model evidence (AIC) but should also successfully simulate participants' sampling behaviour.

To facilitate comparisons between studies, we reanalysed the published findings on the facial attractiveness task using the optimisation method used here. Study 2 as described in Furl et al. (2019) resembled our studies in many ways, including sample size (N = 20), and was therefore chosen to be reanalysed. The only difference between Furl et al. (2019)'s Study 2 and our matchmaker and trustworthiness studies was the decision domain, and thus the instructions given to participants. Henceforth we will refer to this earlier dataset as *facial attractiveness*. Behavioural results from the facial attractiveness dataset were in line with those of matchmaker and trustworthiness; participants oversampled compared to the ideal observer model (t(33.01) = -11.10, $p < .001$) and ended up with lower-ranked options (t(29.78) = 5.84, $p < .001$).

The comparison of participants' sampling behaviour to the computational models was done using fminsearch.m in MATLAB version 2015b (MATLAB, 2015). After the models were generated, the output was exported and further analysed (pairwise t-tests, Bayesian t-tests) as well as visualised in RStudio (RStudioTeam, 2020).

S3 Computational model comparison

A formal model comparison of the log likelihood was performed, however, we are reluctant to draw any conclusions based directly on only the log likelihood as this was uncorrected for the number of parameters (Figure S1a). When comparing the AIC scores, we found no significant difference in all three datasets between the sample reward, biased values, and attractive prior models, using paired *t*-tests Bonferroni corrected for all pairs of models (Figure S1b). All three of these Bayesian models of bias outperformed the cutoff model in every dataset and outperformed the ideal observer model in the facial attractiveness and trustworthiness datasets. Finally, we counted for each model the number of times it came out as the best fitting model for an individual participant. Here, the biased values model appears to be the winner in the facial attractiveness and trustworthiness datasets, but ties with the ideal observer model in the matchmaker dataset (Figure S1c). In the latter, the biased values and ideal observer models differed from the sample reward model by only one participant.

S4 Models correlated with participants' behaviour

The best fitting models identified in the formal model comparison were the sample reward model, the biased values model, and the attractive prior model. In this section, we explore whether the sampling behaviour of the computational models also reliably covaried with participants' sampling data. Code and data for this section are available on the OSF platform[2].

We started with the matchmaker paradigm. Shapiro-Wilk tests of normality indicated that the mean number of samples was normally distributed for all models except the cutoff model (W = 0.90, p = .047). As such, a Spearman's rank correlation was used for the cutoff model, instead of a Pearson correlation. The correlation data validated both the sample reward model and the biased values model (Figure S2).

Next, we performed the same correlation analysis on the trustworthiness data. As above, Shapiro-Wilk tests of normality indicated that the mean number of samples was normally distributed for all models except the cutoff model (W = 0.87, p = .017). As such, a Spearman's rank correlation was used

[2] https://tinyurl.com/yjqy35hd

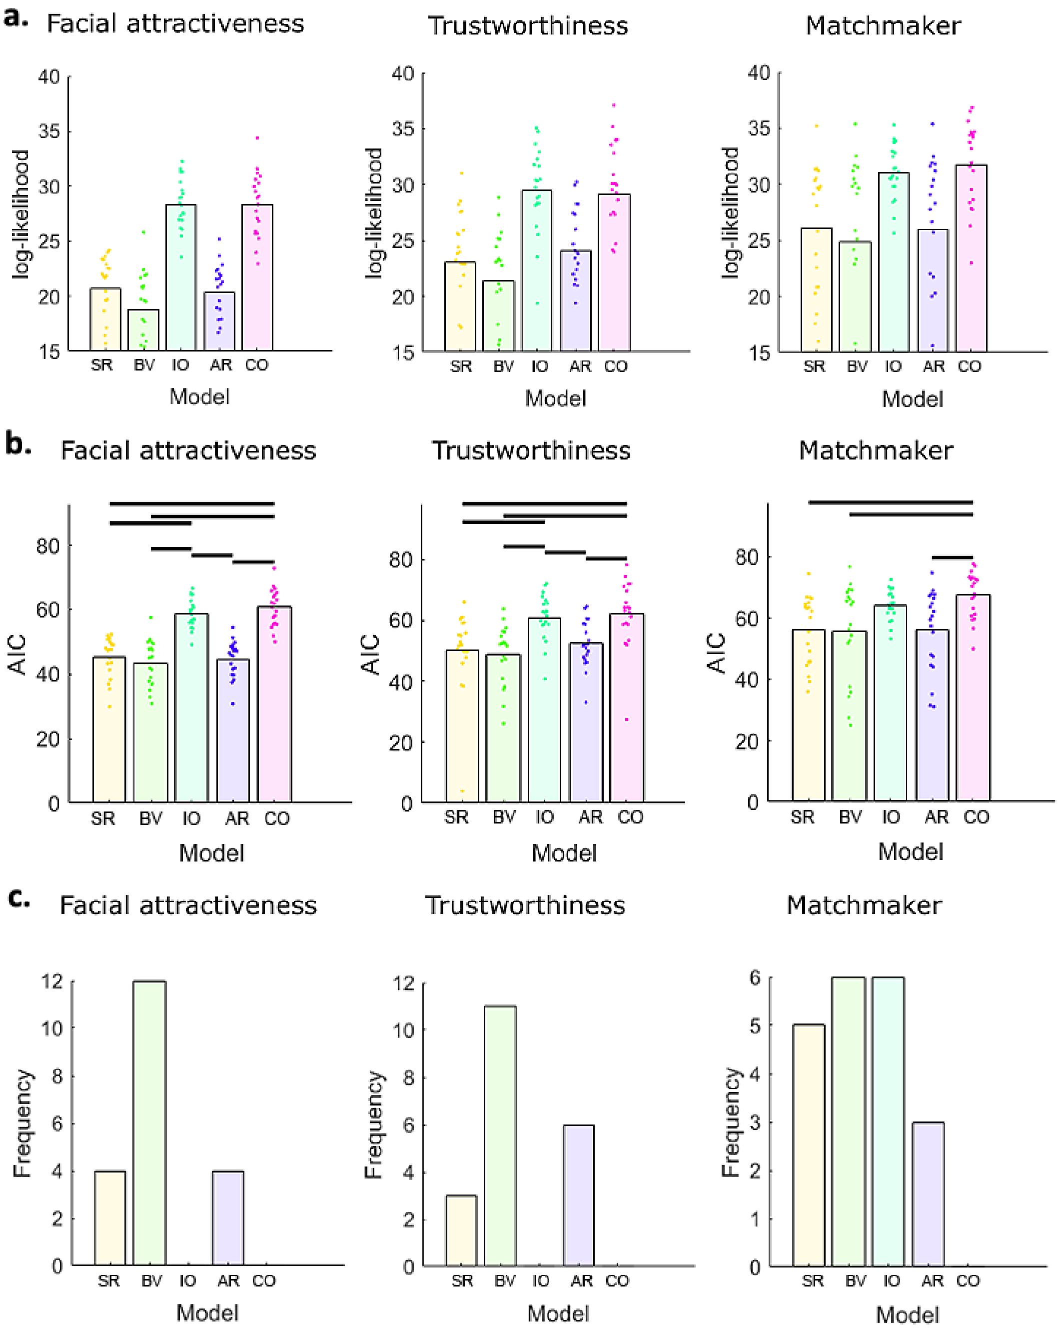

FIGURE S1: Formal model comparison. Top (a): log likelihood for each of the models: sample reward (SR), biased values (BV), ideal observer (IO), attractive prior (AP) and cutoff (CO). Note that a smaller log likelihood indicates a better model fit. Middle (b): AIC scores for each of the same models. Note that a smaller AIC score indicates a better model fit. Black lines denote significant differences between models, calculated using paired *t*-tests Bonferroni corrected for all pairs of models. Bottom (c): Frequency of participants best-fitting to each model. *N* = 20 per dataset.

for the cutoff model, instead of a Pearson correlation. A significant correlation was found for the sample reward model, but not for any of the other models (Figure S3).

Finally, we looked at the facial attractiveness dataset. Shapiro-Wilk tests of normality indicated that the mean number of samples was normally distributed for all models. The correlation data validated the sample reward, attractive prior and biased values models (Figure S4).

In addition to the more traditional frequentist analyses, we present the Bayes Factors for each of the correlations below in Table S2. In accordance with Wagenmakers et al. (2018), we interpret Bayes Factors as follows: $BF_{10} > 100$ = extreme evidence for H1, BF_{10} 30-100 = very strong evidence for H1, BF_{10} 10-30 = strong evidence for H1, BF_{10} 3-10 = moderate evidence for H1, BF_{10} 1-3 = anecdotal evidence for H1 and BF_{10} 1/3-1 = anecdotal evidence for H0. From our Bayes Factor analysis, we can conclude that there is very strong to extreme evidence for a correlation between participants' sampling and that of the sample reward model across all three studies (H1). Furthermore, although there is strong evidence that the sampling rates of the biased values and attractive prior models correlate with participants' sampling rate in the facial attractiveness dataset, the models don't seem to generalise as well to the other two datasets where there is no more than moderate evidence for the biased values model and even anecdotal evidence that there is no correlation (H0) for the attractive prior model.

TABLE S2: Bayes Factors (BF_{10}) for correlations between participants and five different computational models for the mean number of samples before choice, across three different datasets. *** denotes extreme or very strong evidence for H1, ** denotes strong or moderate evidence for H1, * denotes anecdotal evidence for H1.

	Participants		
	Matchmaker	**Trustworthiness**	**Facial attractiveness**
SR	88.536 ±0% ***	222.847 ±0% ***	552.268 ±0% ***
BV	3.366 ±0% **	1.706 ±0% *	13.121 ±0% **
IO	0.557 ±0%	0.489 ±0%	0.870 ±0%
AP	0.493 ±0%	0.572 ±0%	21.115 ±0% **
CO	0.535 ±0%	0.692 ±0%	0.859 ±0%

Based on the results of our correlation analyses, we can conclude that the models identified in the formal model comparison as providing a good fit to participants' data, remain good contenders in predicting participants' sampling behaviour. As such, we cannot draw an unequivocal conclusion about any one model, although the evidence for the sample reward model is very strong.

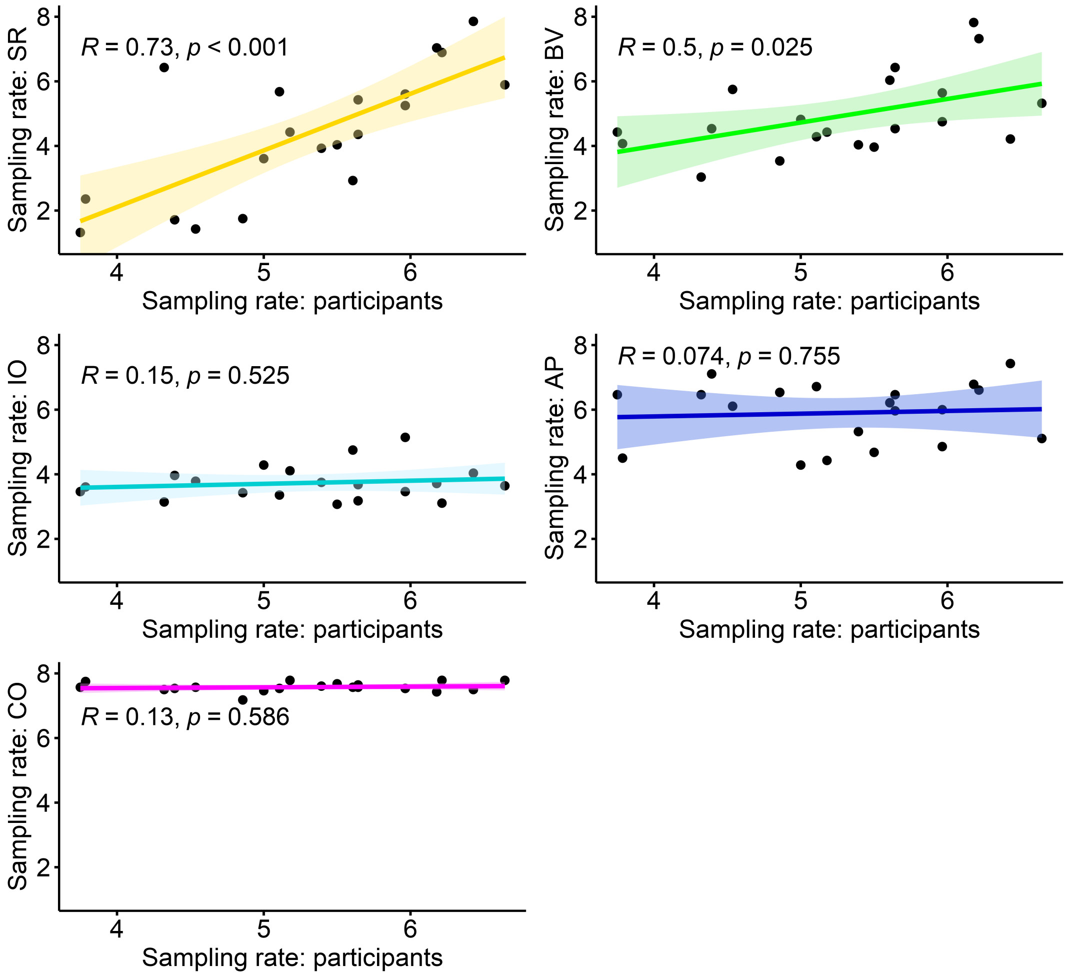

FIGURE S2: *Matchmaker*. Correlations between participants and five different computational models for the mean number of samples before choice. Models are abbreviated as sample reward (SR), biased values (BV), ideal observer (IO), attractive prior (AP) and cutoff (CO). R indicates the correlation coefficient. Shaded areas represent the 95% confidence interval.

S5 Discussion

Our formal model comparison of the sample reward, biased values, ideal observer, attractive prior, and cutoff models show strikingly similar patterns between the three datasets (matchmaker, trustworthiness and facial attractiveness). This indicates that sampling biases across the three datasets may

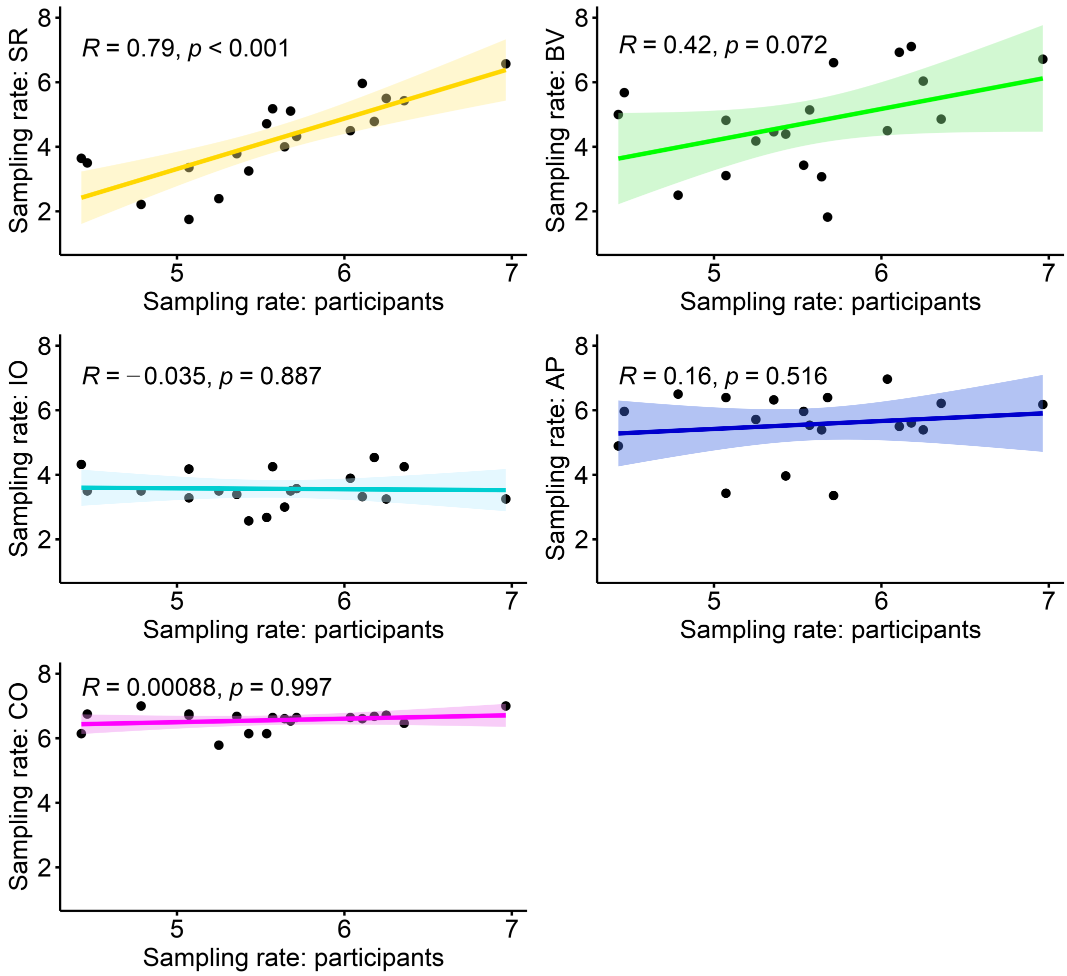

FIGURE S3: *Trustworthiness*. Correlations between participants and five different computational models for the mean number of samples before choice. Models are abbreviated as sample reward (SR), biased values (BV), ideal observer (IO), attractive prior (AP) and cutoff (CO). R indicates the correlation coefficient. Shaded areas represent the 95% confidence interval.

be explained in a similar way. Our results regarding the cutoff model were in agreement with Lee et al. (2005) in that we did not find support for the cutoff model being a suitable model of human decision making on this type of full information optimal stopping task. This conclusion is based on the fact that the cutoff model explained none of the individual subjects' sampling behaviour, it had one of the highest AIC scores, and the model seemed to correlate poorly with participants' mean number of samples across all three datasets. The main contending models based on the AIC (Figure S1b) and the proportion of participants best explained by each model (Figure S1c) seem to be the sample reward model, the biased values model, and the attractive prior model. The biased values model was the 'winning' model in the most

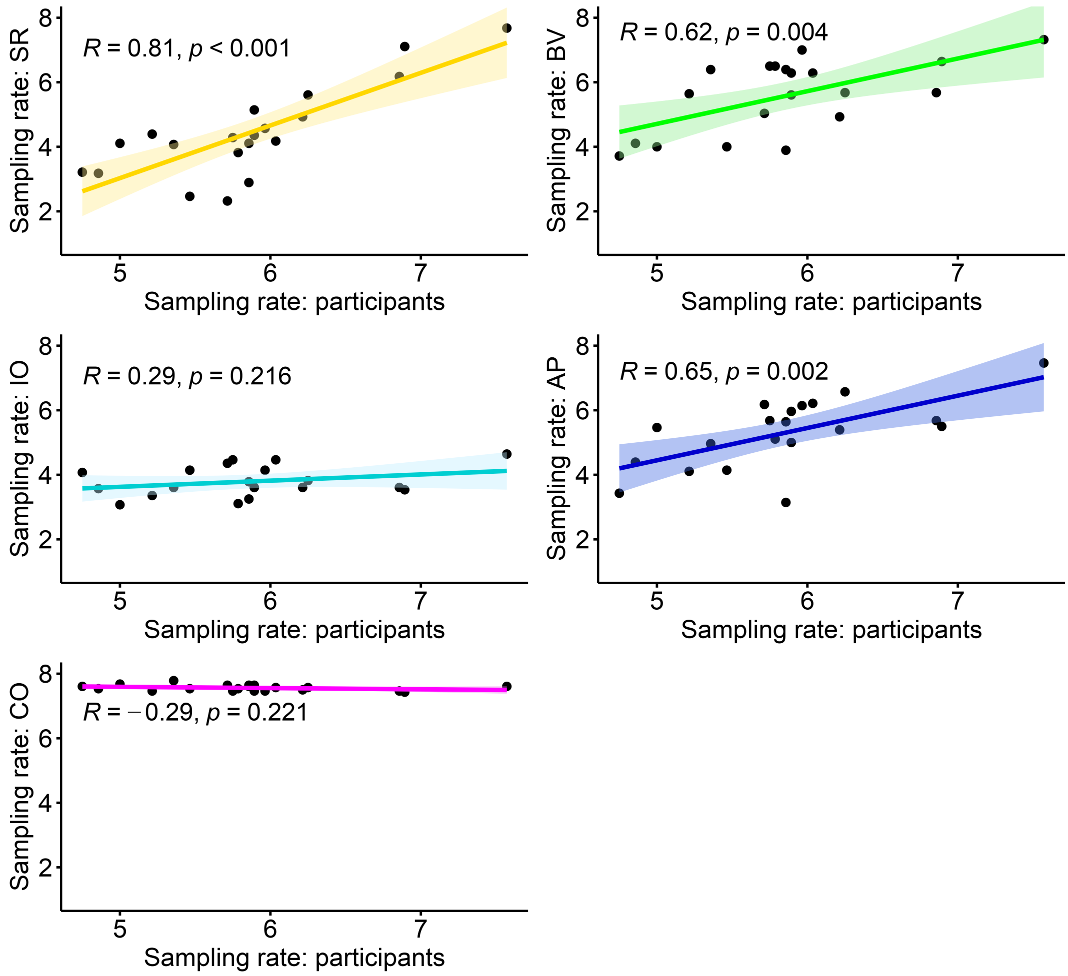

FIGURE S4: *Facial attractiveness.* Correlations between participants and five different computational models for the mean number of samples before choice. Models are abbreviated as sample reward (SR), biased values (BV), ideal observer (IO), attractive prior (AP) and cutoff (CO). R indicates the correlation coefficient. Shaded areas represent the 95% confidence interval.

participants in all three datasets (see Figure S1c), although the ideal observer model tied with the biased values model in the matchmaker dataset. While it might be the case that the ideal observer model was the best-fitting model in these specific participants, the ideal observer model is not likely to contribute to the average oversampling effect observed in the aggregate sample. Looking at the correlations with participants' sampling behaviour, the sample reward, biased values, and attractive prior models presented a good fit to participants' data, with the sample reward model showing the most consistent correlation with human oversampling biases. Indeed, after calculating the Bayes Factor for each of the correlations, very strong evidence across all three datasets was found for the sample reward model correlating with participants. The biased values and attractive prior models did not seem to generalise well to all three datasets, and anecdotal evidence for the null hypothesis was found across the board for the ideal observer and cutoff models.

From these exploratory analyses, we can conclude that the causes of oversampling may be heterogeneous strategies across participants rather than simply overweighting of the most attractive values when making a decision (biased values). The sample reward, biased values and attractive prior models all showed oversampling compared to the ideal observer model, and their sampling rates were not significantly different from participants in either one or both of our studies. To investigate whether oversampling biases are best explained by different models in different participants, we recommend future research should increase the sample size. To demonstrate sampling biases, twenty participants was sufficient both here and in previous studies (Furl et al., 2019). However, comparisons between bias models and the investigation of individual differences could be facilitated by either larger samples or longer experiments with more trials.

Chapter 4

Explaining human sampling rates across different decision domains

Abstract

Undersampling biases are common in the optimal stopping literature, especially for economic full choice problems. Among these kinds of number-based studies, the moments of the distribution of values that generates the options (i.e., the generating distribution) seem to influence participants' sampling rate. However, a recent study reported an oversampling bias on a different kind of optimal stopping task: where participants chose potential romantic partners from images of faces. The authors hypothesised that this oversampling bias might be specific to mate choice. We preregistered this hypothesis and so, here, we test whether sampling rates across different image-based decision-making domains a) reflect different over- or undersampling biases, or b) depend on the moments of the generating distributions (as shown for economic number-based tasks). In two studies (N = 208 and N = 96), we found evidence against the preregistered hypothesis. Participants oversampled to the same degree across domains (compared to a Bayesian ideal observer model), while their sampling rates depended on the generating distribution mean and skewness in a similar way as number-based paradigms. Moreover, optimal model sampling to some extent depended on the the skewness of the generating distribution in a similar way to participants. We conclude that oversampling is not instigated by the mate choice domain and that sampling rate in image-based paradigms, like number-based paradigms, depends on the generating distribution.

1 Introduction

An optimal stopping problem can be defined as a situation in which a decision maker has to choose when to stop searching for more information and take a given action. Optimal stopping problems have long held the fascination of scholars, particularly mathematicians, who were determined to prove that optimal solutions to these kinds of problems exist (for historical reviews, see Ferguson, 1989; Freeman, 1983). Within this paper we focus on a specific and simple version of the so-called full information problem, where the actual values of the options are presented, the distributions that generate the option values (i.e., the generating distributions) are familiar to participants, there is no extrinsic cost-to-sample, there is no recall of rejected options, and decision outcomes provide a reward equal to their value (Abdelaziz & Krichen, 2006; Gilbert & Mosteller, 1966; Guan et al., 2014; Hill, 2009; Lee, 2006; Shu, 2008). These full information decision problems are solved computationally using a backwards induction algorithm (see Chapter 1, Sidebar 1), which predicts the values of future options based on a known distribution that generates the option values (Cardinale et al., 2021; Costa & Averbeck, 2015; Furl et al., 2019; Gilbert & Mosteller, 1966). These models of optimality are programmed by researchers with the mean and variance of the assumed-to-be-normal generating distribution (or the prior of this distribution), which the researchers assume the participants are using (Costa & Averbeck, 2015; Gilbert & Mosteller, 1966). Other types of optimal stopping problems exist that require somewhat different computational solutions, but those are outside the scope of this paper (e.g., Goldstein et al., 2020; Van der Leer et al., 2015; Zwick et al., 2003).

Previous research has found evidence that participants commonly undersample compared to optimality in our focus case of full information problems (Baumann et al., 2020; Cardinale et al., 2021; Costa & Averbeck, 2015). This undersampling bias has also been found in multiple closely related optimal stopping problems that go beyond our focus, including the classic secretary task (Bearden et al., 2006; Seale & Rapoport, 1997), numerical optimal stopping tasks (Guan & Lee, 2018; Guan et al., 2014; Kahan et al., 1967; Shapira & Venezia, 1981), and even the beads task (Furl & Averbeck, 2011; Hauser et al., 2017; Hauser et al., 2018; Van der Leer et al., 2015). In addition to undersampling biases, human sampling rates can be affected by other factors such as sequence length (Cardinale et al., 2021; Costa & Averbeck, 2015; Goldstein et al., 2020), cost-to-sample (Costa & Averbeck, 2015;

Zwick et al., 2003), or the moments of the distribution that generates the option values (i.e., many high/low value options; Baumann et al., 2020; Guan & Lee, 2018; Guan et al., 2014; Guan & Stokes, 2020). The latter factor - moments of the generating distribution - is a focus of the current paper.

However, a recent study reported an oversampling bias in a different decision-making domain within the same full information modelling framework. In a mate choice decision scenario, participants searched for the most attractive date from a series of faces (Furl et al., 2019). Furl et al. (2019) hypothesised that the oversampling bias on this non-economic, image-based task might be specific to the mate choice decision-making domain. Their hypothesis was based on behavioural ecology research which suggests that animals use high thresholds for mate choice (Backwell & Passmore, 1996; Ivy & Sakaluk, 2007; Milinski & Bakker, 1992). Furthermore, Furl et al. (2019) reported that a computational model that incorporated such a high-threshold bias best described participants' sampling behaviour.

Here, we continue to examine influences on human sampling rate in image-based optimal stopping tasks with two main hypotheses. The first is pre-registered and based on Furl et al. (2019)'s proposal that mate choice selectively leads to oversampling. This account predicts that image-based domains other than facial attractiveness should not lead to oversampling biases. The second hypothesis is based on previous studies of the number-based full information task which show that a more positively skewed generating distribution can increase sampling rate (Baumann et al., 2020). To date, this hypothesis has been tested in number-based domains only, and has yet to be tested in image-based domains. For the two purposes outlined above, we have chosen three image-based decision-making domains: faces (replication), food, and holiday destinations.

2 Materials and methods

We conducted two studies aimed at convergent results; one online (Study 1) and one in a classroom setting (Study 2). The data analysis plan for our online study was preregistered before data collection, and is openly available on the AsPredicted pre-registration website.[1] Our classroom study was not separately pre-registered but followed the same data analysis protocol

[1]https://aspredicted.org/sr5fv.pdf

as pre-registered for Study 1. Methods for Study 1 and Study 2 were nearly identical, as outlined below.

2.1 Study 1

Participants

Both studies were approved by Royal Holloway, University of London's Ethics Board. Informed consent was obtained from all participants before the start of the study, in accordance with the Declaration of Helsinki. For our first study, 225 participants were recruited through the online recruitment service Prolific (Prolific, 2014). Two participant prerequisites were set, the first being age (between 18 and 35), as this roughly matched the age range of the faces shown in the study. The second prerequisite was nationality (either United Kingdom (UK), Ireland, United States (US), Canada, Australia, or New Zealand), which was set under the assumption that participants with these nationalities would have a good command of the English language, and would therefore be able to sufficiently understand the instructions and the informed consent form. Each participant was randomly assigned to one of three conditions (N = 75 each), with each condition corresponding to a different decision domain. Participants received a flat fee as compensation for completing the study, with the entire study lasting about 15 minutes.

Paradigm

Gorilla Experiment Builder (Anwyl-Irvine et al., 2020) was used to create and host our studies. The paradigm for all three domains (faces, food and holiday destinations) was very similar, and inspired by the methods used across the three studies described in Furl et al. (2019). The paradigm consisted of two phases; a rating phase and a sequence phase (i.e., the optimal stopping task). Before commencing the study, participants in the faces domain were asked to choose whether they would like to rate (and date) males or females. Based on their answers, each was shown either male or female faces throughout the study.

In phase one of the study, participants rated 180 images in total (90 unique images, all rated twice) using a slider scale ranging from very unattractive (value 1) to very attractive (value 100). Consistent with previous studies of full information problems (Cardinale et al., 2021; Costa &

Averbeck, 2015), the prior of the option-generating distribution was configured with the mean and variance of the generating distribution. In our case, as in Furl et al. (2019), this means the distribution of subjective values (attractiveness ratings). Using personalised ratings ensures that the likelihood of an image being chosen is not influenced by individual differences in attractiveness preferences (Furl et al., 2019), as the same prior distribution of values was available for learning to both agents, that is, the participants and the optimality model against which we compare the participants (see Section 2.3). Sliders were made invisible until first click to reduce slider biases (Matejka et al., 2016), and the slider's current selected value was shown for increased precision. A progress bar was shown at the bottom of the screen to visualise participants' progression. An attention check was included in phase one to compensate for the unsupervised nature of online data collection (see Supplementary Materials, Section S1). Final attractiveness ratings were computed from the mean of the two ratings, which previous work has found to be sufficient for detecting oversampling on the facial attractiveness paradigm (Furl et al., 2019, Study 3) and which shortened the duration of our study to suit online presentation. Internal consistency between the two ratings, measured using Cronbach's alpha, was acceptable (Taber, 2018), confirming that participants were consistent in their ratings of images (female faces: α = .848, male faces: α = .882, food: α = .954, holiday destinations: α = .926).

For the faces domain, 90 faces were randomly selected from a larger set of 426 images, the same set used in Study 2 of Furl et al. (2019). The set of 90 food images was randomly selected from a larger set of 1314 images (Blechert et al., 2019). The image numbers corresponding to the food images that were used in this study can be found in the Supplementary Materials (Section S2). The set of holiday destination images was randomly selected from a royalty-free image database (*www.shutterstock.com*). Search terms that were used included, for example, 'holiday destination', 'holiday', 'travel destination', 'travelling', and 'European city'. Stimulus dimensions of the three stimulus sets were kept as homogeneous as possible. For example, all images were cropped to the same size (1200 pixels) and the same shape (square). Other stimulus dimensions such as hue and saturation were not further controlled for, as differences can be expected both within and between domains.

In the second phase, participants were shown six sequences of eight images each. Images were randomly sampled from the entire distribution of images that had been rated in phase one. Participants attempted to choose the most attractive option from the sequence as they could, with the restriction that they could not return to a previously rejected option. The number of options remaining was shown at the top of the screen, and the rejected options were shown at the bottom of the screen. When a participant made a choice, they had to advance through a series of grey squares that replaced the remaining images. This ensured that participants could not finish the experiment early by choosing an early option. Adding grey squares does not alter participants' sampling behaviour: Furl et al. (2019) found the same results on the facial attractiveness paradigm with the implementation of grey squares (Studies 2 and 3) and without (Study 1). The entire study was self-paced - participants advanced by using their mouse to click on the buttons on the screen. If the last option in the sequence was reached, that option became their choice by default. After finishing a sequence, participants were directed to a feedback screen displaying the participant's chosen image, and the text: "Here is your [new date / next meal / next holiday destination]! How rewarding is your choice?". Participants responded to this question using a slider scale ranging from not rewarding (value 1) to very rewarding (value 100). The feedback screen was included to provide feedback about the quality of the participants' choice by asking them to reflect upon its reward value before moving onto the next sequence, and responses were not further analysed. Next, participants were directed to a screen asking them: "Ready for the next sequence?". Participants responded by clicking a button saying: "I'm ready!".

The two key dependent variables of interest are the position of the chosen image in the sequence (i.e., number of samples), and the rank of the chosen image (out of the images in the sequence). Both variables are a mean value over six sequences for each participant.

2.2 Study 2

A second study was conducted in a laboratory setting to replicate the results of Study 1 (which was conducted online) and thus bolster our findings. Opportunity sampling was used to recruit 96 participants during an Open Day at Royal Holloway, University of London. This sample size was sufficiently large, as a power analysis based on the outcomes of Study 1 indicated that for Study 2, a total sample size of 70 participants was sufficient for 95%

power.[2] Participants were randomly allocated to one of three domains, with final numbers being 32 in faces, 28 in food, and 36 in holiday destinations. Participants did not receive any monetary compensation for their participation.

Paradigm

Study 2 used a shortened, but otherwise identical version of the paradigm described in Section 2.1. The reason why it was shortened was because of time constraints related to the recruitment format (during a University Open Day). As such, the two key differences between Study 2 and Study 1 are 1) participants in Study 2 rated every image only once, and 2) the attention check was removed in Study 2 as the study was conducted in a more controlled setting. The shortened version of the paradigm might introduce more noise in the data, which could reduce our ability to detect a result. Despite this, we found no differences in the results due to the shortened format.

2.3 Optimal model

Participants' sampling behaviour was compared to a Bayesian ideal observer model (Costa & Averbeck, 2015), where performance is Bayesian optimal and the cost-to-sample parameter was fixed to zero. This model has previously been used by e.g., Costa and Averbeck (2015), Furl et al. (2019) and Cardinale et al. (2021), and is the same as the model of Gilbert and Mosteller (1966) in that both assume that options are sampled from a normal option generating distribution with known mean and variance, and both use a backwards induction algorithm to compute the value of sampling again, which is compared to the value of the current option. The Bayesian optimality model enhances the original Gilbert and Mosteller model by adding to it 1) a generating distribution that is initialised with a prior distribution which is then updated after each new sample using Bayes' rule, 2) a cost-to-sample parameter (here set to zero), and 3) functionality for the researcher to apply any arbitrary reward function to the choice outcomes. Mathematically, the model is based on a discrete time Markov decision process with continuous states. Theoretically, at each position in the sequence, the optimal model computes

[2] The smallest effect size measured in Study 1 was $f = 0.484$ (for the comparison between the conditions faces, food, and holiday destinations), which for an "ANOVA fixed effects, special, main effects and interactions" statistical test, which we used to analyse Study 1 and as such is suited for our analysis of Study 2, leads to a sample size of 70 participants for 95.3% power, as calculated in G*Power (Faul et al., 2009).

the respective values for choosing the option and declining the option, and chooses the one with the highest value. To calculate the value of either taking or declining an option in a sequence the model computes the action value Q as:

$$\begin{aligned} &Q_t\left(s_t, a = take\right) = r_t\left(s_t, a\right) \\ &Q_t\left(s_t, a = decline\right) = \int_s p_t\left(j \mid s_t, a\right) u_{t+1}(j) dj \end{aligned} \tag{4.1}$$

The key computations for the optimal model, as seen in equation 4.1, are utility (equation 4.2) and reward values (equation 4.3). The model uses backwards induction to derive utilities that could result from further sampling (equation 4.4).

$$u_t\left(s_t\right) = \max_{a \in A_{s_t}} \left\{ r_t\left(s_t, a\right) + \int_s p_t\left(j \mid s_t, a\right) u_{t+1}(j) dj \right\} \tag{4.2}$$

The utility u of the state s at sample t is the value of the best action a, which depends on reward value r, the cost-to-sample C_s, and the probabilities of outcomes j of subsequent states, weighted by their utilities. i represents each option in a sequence.

$$\begin{aligned} &r_t\left(s_t, a = accept\right) = \sum_{i=1}^{N} p\left(rank = i\right) * R\left(i + (h-1)\right) \\ &r_t\left(s_t, a = decline\right) = C_s \end{aligned} \tag{4.3}$$

Our optimal model adds to the Gilbert and Mosteller model a function R, which maps the rank of each option to the amount of reward gained when choosing an option of that rank. We assumed that participants followed our instructions and tried to choose the option with the highest subjective value possible. The corresponding model, if it followed these instructions, would therefore gain a reward commensurate with the subjective value (rating) of the chosen option. That is, we assigned to R(1) the rating of the highest ranked option, R(2) the rating of the second highest ranked item, and so on. This reward function resembles the classic Gilbert and Mosteller model, which also attempts to maximise the option value of its choices. h represents the relative rank of the current option. When considering final sequence position N, the model computes final utilities as:

$$u_N(s_N) = r(s_N) \text{ for all } s_N \in N \tag{4.4}$$

and working backwards from N, we use equation 4.2 to compute utilities at every sequence position t.

The value for declining an option can be considered the choice threshold, as no option is chosen unless the value for choosing an option exceeds the value for declining an option. The choice threshold is dynamic, and can change depending on the position in the sequence. The model received as input for each participant the values of the sequence options as presented to the participant in phase two, with each sequence value comprising the participant's individual rating of the option. To approximate normality, ratings were log transformed for each participant before being put into the model. Input and parameter settings for the optimal model described here apply to all analyses in this paper.

2.4 Data analysis

The comparison of participants' sampling behaviour to the optimal model was done using MATLAB version 2015b (MATLAB, 2015). Statistical tests were performed using RStudio (RStudioTeam, 2020). For all analyses, a p value of $< .05$ was considered significant. Additionally, to allow evidence for the null hypothesis to be quantified, we show the Bayes factors for mean number of samples and mean rank as well. Bayesian t-tests were calculated using the BayesFactor package (Morey & Rouder, 2018), within the R environment. We follow guidelines provided by Wagenmakers et al. (2018) to interpret Bayes factors, with $BF_{10} > 100$ being interpreted as extreme or decisive evidence for the alternative hypothesis, and $BF_{10} < .01$ being decisive evidence in favour of the null model (no differences between means).

3 Results

3.1 Study 1

After the removal of any outliers (see Supplementary Materials, Section S3), the final number of participants in each domain was 68 for faces, 72 for food, and 68 for holiday destinations (for demographic statistics, see Table S1 in the Supplementary Materials). In the facial attractiveness domain, the majority of participants chose to rate faces of the opposite sex (89.7%).

Because we were interested in testing the hypothesis proposed by Furl et al. (2019) that oversampling bias is specific to the mate choice domain, we pre-registered the hypothesis that there would be a significant domain*agent interaction. We therefore implemented a 3x2 factorial ANOVA to compare the differential effects of our two agents (participants and model) across the three domains. We found that the domain*agent interactions for the mean number of samples and the mean rank of the chosen option did not reach significance (Table 4.1), so there was no evidence of a difference in sampling bias between domains. This is confirmed by the Bayes factor analysis, which showed that there was no evidence that the full model (domain + agent + domain*agent) was better than just the domain + agent model (BF_{10} = 0.260). In fact, there was extreme evidence for the domain + agent model ($BF_{10} = 6.790e^{49} \pm 1.08\%$). This Bayesian analysis, therefore, provides positive evidence for the absence of our pre-registered domain*agent interaction.

Because the domain*agent interaction effect was not significant, this meant that, on average, the sampling rates of the two agents (participants and model) varied in the same way across domains. Indeed, when looking at sampling biases, we found evidence that despite variations in sampling rate for both agents across the domains (Figure 4.1), participants oversampled in each of our three domains when tested separately (Table 4.2), and achieved lower ranks than the optimal model (Figure 4.2, Table 4.2). Bayes factor t-tests supported this finding, showing extreme evidence for a difference between participants and the optimal model for the mean number of samples as well as the mean rank for each of the three domains (Table 4.3). Collapsing over agents, agents on average sampled more and achieved higher ranks in the faces domain than in either of the other two domains. Furthermore, agents on average sampled more in the food domain than in the holiday destinations domain (Table 4.4).

We also tested for the effect of self-reported participant sex on the two dependent variables, mean number of samples and mean rank of the chosen option, but did not find significant results (F(1, 73) = 1.279, p = .262 and F(1, 73) = 1.814, p = .182, respectively).

TABLE 4.1: 3x2 factorial ANOVA describing the main effects and interaction effects for the mean number of samples and the mean rank of the chosen option, in both Study 1 and Study 2. Degrees of freedom is abbreviated as DF.

	Study 1			**Study 2**		
	DF	F	p	DF	F	p
Number of samples						
Agent	(1, 406)	240.75	<.001	(1, 185)	159.27	<.001
Domain	(2, 406)	46.83	<.001	(2, 185)	27.81	<.001
Agent*Domain	(2, 406)	1.77	0.171	(2, 185)	3.70	0.023
Rank						
Agent	(1, 408)	232.53	<.001	(1, 186)	18.57	<.001
Domain	(2, 408)	7.43	<.001	(2, 186)	1.76	0.175
Agent*Domain	(2, 408)	0.97	0.382	(2, 186)	6.94	0.001

TABLE 4.2: Post hoc Friedman's tests (Bonferroni corrected for the three domains) to test for differences between agents in each individual domain, in both Study 1 and Study 2.

	Study 1	**Study 2**
Number of samples		
Faces	$< .001$	$< .001$
Food	$< .001$	$< .001$
Holidays	$< .001$	$< .001$
Rank		
Faces	$< .001$	$< .001$
Food	$< .001$	.450
Holidays	$< .001$	.304

3.2 Study 2

Removed outliers and demographic statistics for Study 2 can be found in the Supplementary Materials (Section S4 and Table S2). In the facial attractiveness domain, 78.1% of participants chose to rate faces of the opposite sex.

Unlike Study 1, Study 2 achieved a significant interaction between agent and domain for both the mean number of samples and the mean rank of the chosen image (Table 4.1). Upon visual inspection of Figure 4.1, we hypothesised that these interactions arose as a result of the magnitude of the

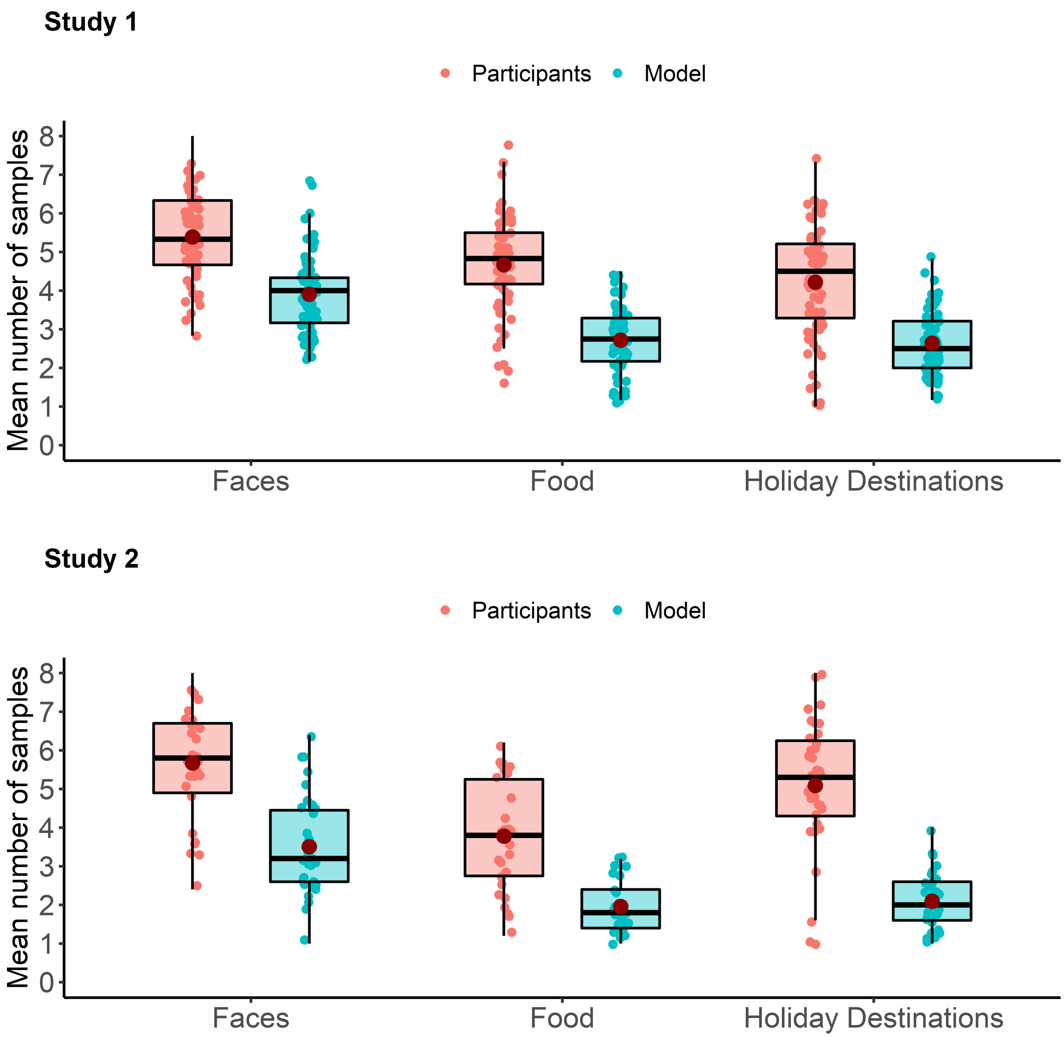

FIGURE 4.1: Box plots and raw jittered data points for the mean number of samples for participants versus the optimal model, grouped by domain. The red dots represent the mean, horizontal black lines represent the median, boxes show the 25% and 75% quantiles, and the whiskers represent the 95% confidence intervals.

oversampling bias varying per condition. That is, the difference between the mean sampling rate of participants and the model is 2.16 options in the face condition, 1.82 options in the food condition, and 2.99 options in the holiday destinations condition. The difference between agents is significant in all conditions (see post hoc results in Table 4.2). The fact that we found that participants varied in sampling rate from domain to domain is in line with our findings of Study 1. Nevertheless, oversampling is generally maintained because the model most of the time adjusts from domain to domain

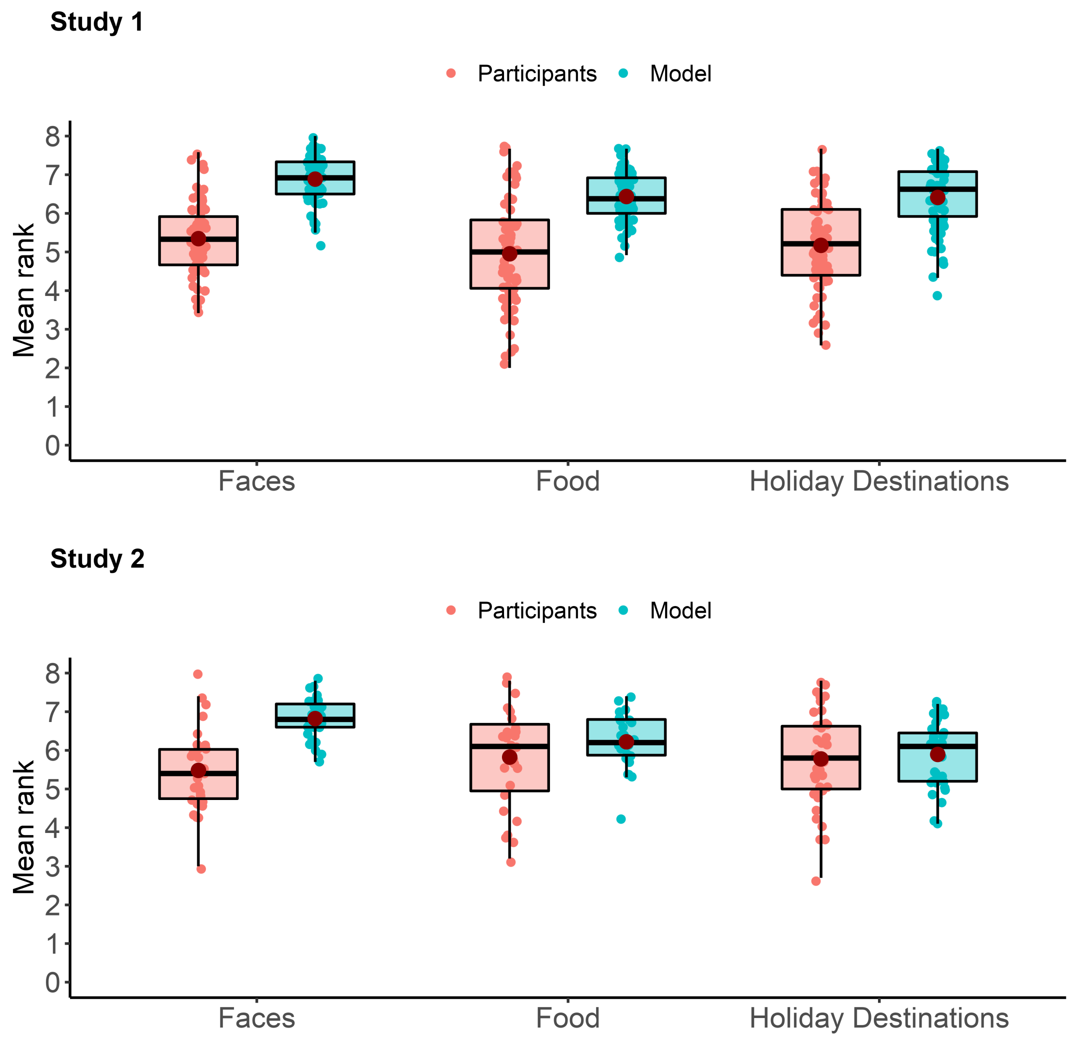

FIGURE 4.2: Box plots and raw jittered data points for the mean rank of the chosen option for participants versus the optimal model, grouped by domain. The red dots represent the mean, horizontal black lines represent the median, boxes show the 25% and 75% quantiles, and the whiskers represent the 95% confidence intervals.

the same as participants. Indeed, we found evidence that participants oversampled relative to the optimal model in each of our three domains (Tables 4.2 and 4.3). Collapsing over agents, agents on average sampled more in the faces domain than in the other two domains. Furthermore, agents on average sampled more in the holiday destinations domain than in the food domain (Figure 4.1, Table 4.4). We did not find any significant differences in the mean rank between the three domains (Figure 4.2, Table 4.4).

TABLE 4.3: Bayes factor (BF_{10}) describing the difference between agents for the mean number of samples and the mean rank of the chosen option, for each of the three domains.

	Study 1	Study 2
Number of samples		
Faces	$4.006e^{9}$ ±0%	$1.408e^{5}$ ±0%
Food	$1.790e^{14}$ ±0%	$8.753e^{4}$ ±0%
Holidays	$8.345e^{7}$ ±0%	$5.395e^{7}$ ±0%
Rank		
Faces	$6.098e^{16}$ ±0%	$7.273e^{4}$ ±0%
Food	$1.478e^{9}$ ±0%	0.460 ±0.03%
Holidays	$1.841e^{8}$ ±0%	0.201 ±0%

TABLE 4.4: Post hoc pairwise *t*-tests describing the main effects of domain (averaged over agents) on the mean number of samples and the mean rank of the chosen option, in both Study 1 and Study 2. *p* values are corrected using Fisher's Least Significant Differences.

	Study 1			Study 2		
	Faces	Food	Holidays	Faces	Food	Holidays
Number of samples						
Faces						
Food	<.001			<.001		
Holidays	<.001	.047		<.001	.002	
Rank						
Faces						
Food	<.001			.479		
Holidays	<.001	.394		.065	.289	

4 Generating distribution moments can predict the number of samples

The results of our two studies provide convergent evidence that participants oversample across all three domains, indicating that qualitatively different biases do not explain sampling rates in different domains, as we pre-registered and as hypothesised by Furl et al. (2019). Yet it remains possible that the moments of the generating distribution might explain the variations in sampling rate across these different image-based domains, as they do for economic number-based tasks (Baumann et al., 2020; Guan & Lee, 2018). Presently it is unknown whether the moments of the generating distributions for different domains can explain the differential sampling rates among those domains and it remains to be tested whether the same effects of generating distribution moments seen on economic number-based tasks also hold for the facial attractiveness or other image-based domains. These questions we further examine in the following exploratory analyses.

Figure 4.3 shows the kernel densities of the generating distributions in Study 1 and Study 2. It should be noted that within the facial attractiveness domain, there were essentially two sets of mutually exclusive images: male and female faces, which are here plotted separately. Visual inspection of the density plots potentially suggests marked differences among the four domains in the mean, variance, skewness and kurtosis of the distributions of attractiveness ratings. For example, both studies show a pattern where facial attractiveness ratings appear to have lower means and to be more positively skewed. By contrast, the food domain appears to have a higher mean and to be more negatively skewed. Both studies appear to show the same patterns of distribution shapes, indicating that these distributions are not entirely idiosyncratic from participant to participant, but systematically vary on average over participants too. Additionally, Figure 4.3 shows that our manipulation of the stimulus domain is effectively also an experimental manipulation of the shape of the prior distribution.

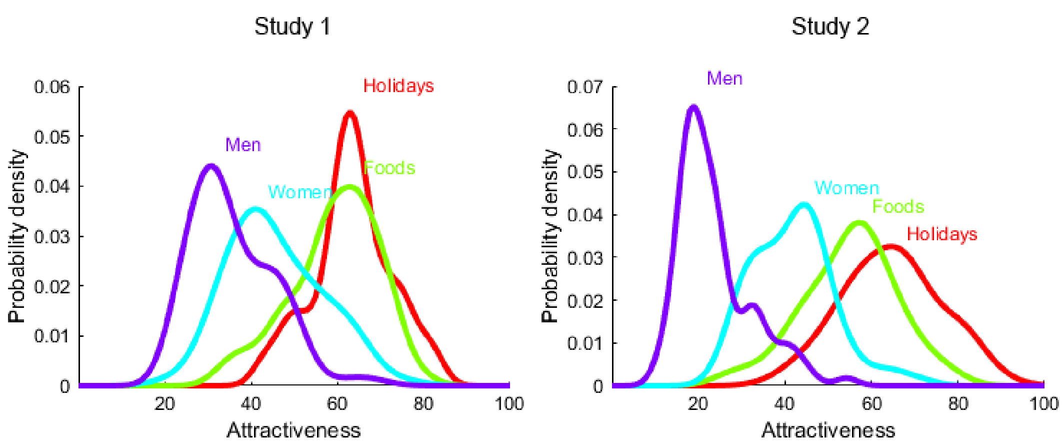

FIGURE 4.3: Density plots for both Study 1 and Study 2 visualising the generating distribution of option values for each domain, with male and female faces plotted separately.

If generating distribution moments affect sampling rates in the different image-based domains in the same way as in economic number-based tasks utilised by for example Guan and Lee (2018) and Baumann et al. (2020), then we would expect greater sampling for more positively skewed (i.e., scarce) environments, like the faces domain. First, we plotted for each participant the mean (Figure 4.4a), variance (Figure 4.4b), skewness (Figure 4.4c), and kurtosis (Figure 4.4d) of their rating distribution. Parallel independent findings were obtained for both Study 1 and Study 2. Participants and data points that were identified as outliers (see Sections S3 and S4 in the Supplementary Materials) remained excluded from the analysis. Additionally, we observed two extreme outliers for kurtosis (65.19 and 26.66), so these two participants were removed from the analysis as well (one in faces and one in holiday destinations). From Figure 4.4 we can observe that the domains male, female, food, and holiday destinations have robustly increasing mean values and decreasing skewness values, consistent with our visual interpretation of the densities in Figure 4.3.

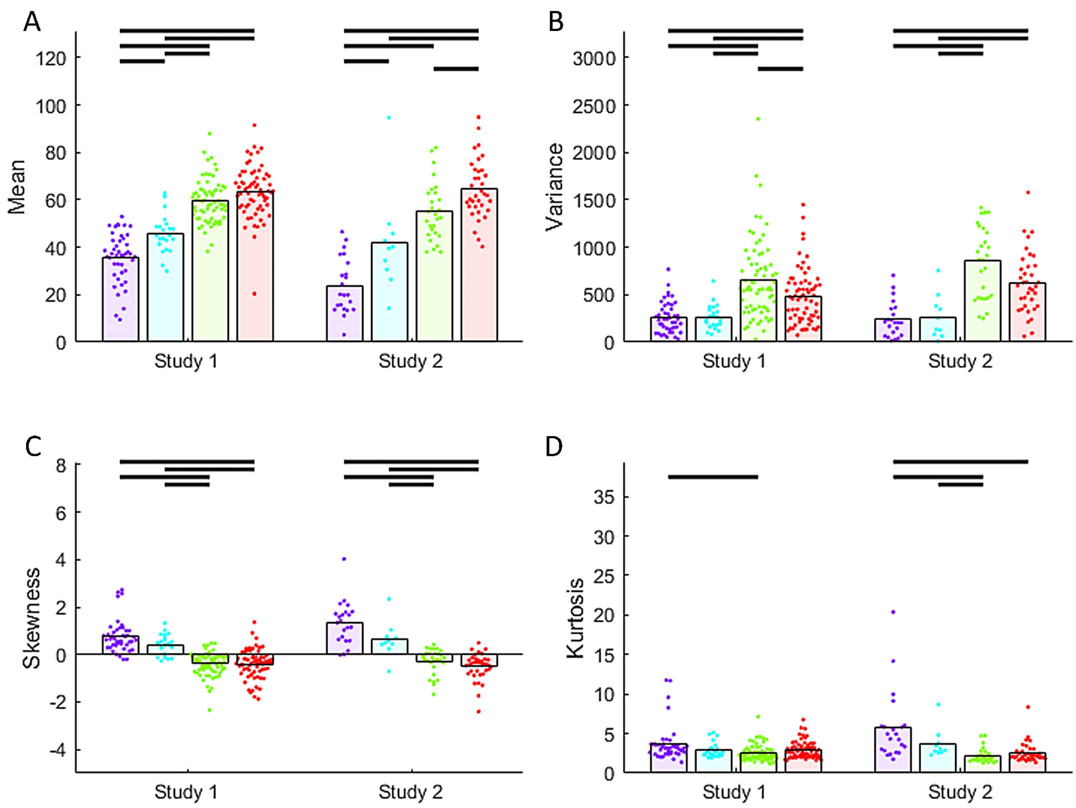

FIGURE 4.4: Generating distribution moments plotted for each of the four domains male faces (purple), female faces (cyan), food (green), and holiday destinations (red), for Study 1 and Study 2. Black horizontal lines denote significant differences between domains at $p < .05$ Bonferroni-corrected for all pairs.

Next, we investigated whether generating distribution moments could predict sampling behaviour using a multiple linear regression model (Figure

4.5). We note that there was multicollinearity between the mean and skewness values of the distributions in both Study 1 (VIF = 3.56 and VIF = 3.83 respectively, r = -0.84) and Study 2 (VIF = 7.84 and VIF = 9.72 respectively, r = -0.92). This means that it is impossible to empirically untangle whether mean and skewness have separate effects on sampling rates. However, as multicollinearity does not affect the predictions, precision of the predictions, and the goodness-of-fit statistics, we continued with a single predictor regression analysis as detailed below. Reported p values should be interpreted with caution.

We found that the mean value was able to predict sampling rate in both Study 1 (β = -0.030, $t(201)$ = -4.926, $p < .001$) and Study 2 (β = -0.024, $t(94)$ = -2.923, p = .004). In Study 1, the mean of the generating distribution explained 10.8% of variation in participants' mean number of samples, while in Study 2 the mean value explained 8.3% of variance. We also found that the skewness of the generating distribution predicted sampling rate in both Study 1 (β = 0.523, $t(201)$ = 4.332, $p < .001$) and Study 2 (β = 0.542, $t(94)$ = 3.191, p = .002). In Study 1, the skewness explained 8.5% of variation in participants' mean number of samples, while in Study 2, the skewness explained 9.8% of variance. These effects are in line with previous findings from full information problems using number-based tasks showing greater sampling in scarce environments (i.e., option generating distributions with lower means and/or more positive skew) (Baumann et al., 2020; Guan & Lee, 2018; Guan et al., 2014). We did not find any significant results for the moments variance (Study 1: β = -0.0005, $t(201)$ = -1.668, p = .097; Study 2: β = -0.0008, $t(94)$ = -1.761, p = .081) and kurtosis (Study 1: β = -0.054, $t(201)$ = -0.818, p = .414; Study 2: β = 0.068, $t(94)$ = 1.029, p = .306).

As we observed participants and the optimal model to both sample more in the faces domain than the other two domains, it is possible that the moments of the generating distribution affected the model in the same way as they did the participants. Single predictor regression analysis of the model's sampling rate and generating distribution moments showed similar effects to participants (Table 4.5, Figure 4.6). Our finding that both participants and the optimal model are sensitive to the moments of the generating distribution could explain the heightened sampling rate in the faces domain, compared to the food and holiday destination domains.

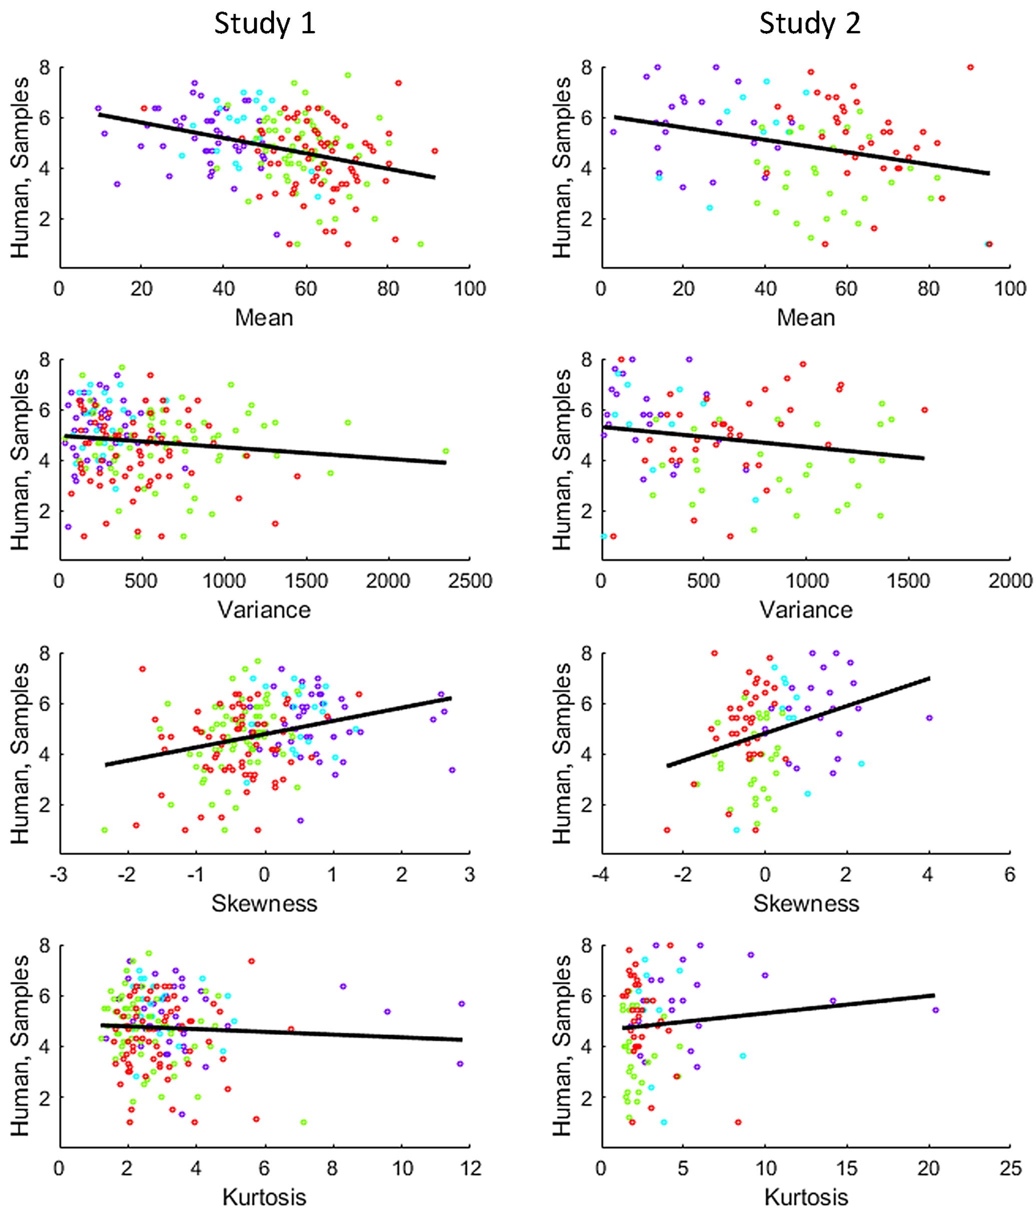

FIGURE 4.5: Scatterplots of generating distribution moments and the mean number of samples for each participant, for Study 1 and Study 2, separated by domain: male faces (purple), female faces (cyan), food (green), and holiday destinations (red). The black line is the regression line.

5 Discussion

The two studies described here addressed our pre-registered hypothesis, derived from Furl et al. (2019), that mate choice is a special decision-making domain that provokes an oversampling bias on optimal stopping tasks, and that this oversampling bias cannot be observed in other domains (e.g., food, holiday destinations). However, contrary to our a priori expectations, we

TABLE 4.5: Single predictor regression analysis of the optimal model's sampling rate and generating distribution moments. *** $p < .001$, ** $p < .01$, * $p < .05$.

Study	Moment	Coefficients	% variance explained
Study 1	Mean	β = -0.014, $t(201)$ = -3.557 ***	5.9%
	Variance	β = -0.001, $t(201)$ = -8.153 ***	24.9%
	Skewness	β = 0.360, $t(201)$ = 4.793 ***	10.3%
	Kurtosis	β = 0.129, $t(201)$ = 3.178 **	4.8%
Study 2	Mean	β = -0.012, $t(94)$ = -3.004 **	8.8%
	Variance	β = -0.001, $t(94)$ = -4.131 ***	15.4%
	Skewness	β = 0.367, $t(94)$ = 4.623 ***	18.5%
	Kurtosis	β = 0.099, $t(94)$ = 3.194 **	9.8%

found in both studies that oversampling generalised to all three image-based decision-making domains. Furthermore, our results were consistent with our second hypothesis, by which participants might oversample across many diverse image-based domains.

Specifically, we found that while different domains did not lead to qualitatively different biases, sampling rates in these domains were increased for positively skewed (i.e., scarce) option generating distributions, consistent with other work using number-based full information tasks (Baumann et al., 2020). What led to this conclusion was the observation that there were modulations in sampling rate for both agents (participants and model) from domain to domain. For example, both participants and the optimal model sampled more in the faces domain (compared to the food and holiday destination domains), while the faces domain also had the lowest mean and most positively skewed generating distribution (Figures 4.1 and 4.3). As such, we suggest that the mean and skewness of the generating distribution could statistically explain sampling behaviour. This is further supported by our finding that the optimal model's sampling rate also correlated with the moments in a similar way to participants. In other words, the moments of the generating distribution lead to the same domain-related variations in sampling rates for both participants and the model.

One of the novel aspects of our study is that we go beyond artificial experimental environments and show that natural image domains for realistic decision problems have variations in distribution shape that can affect

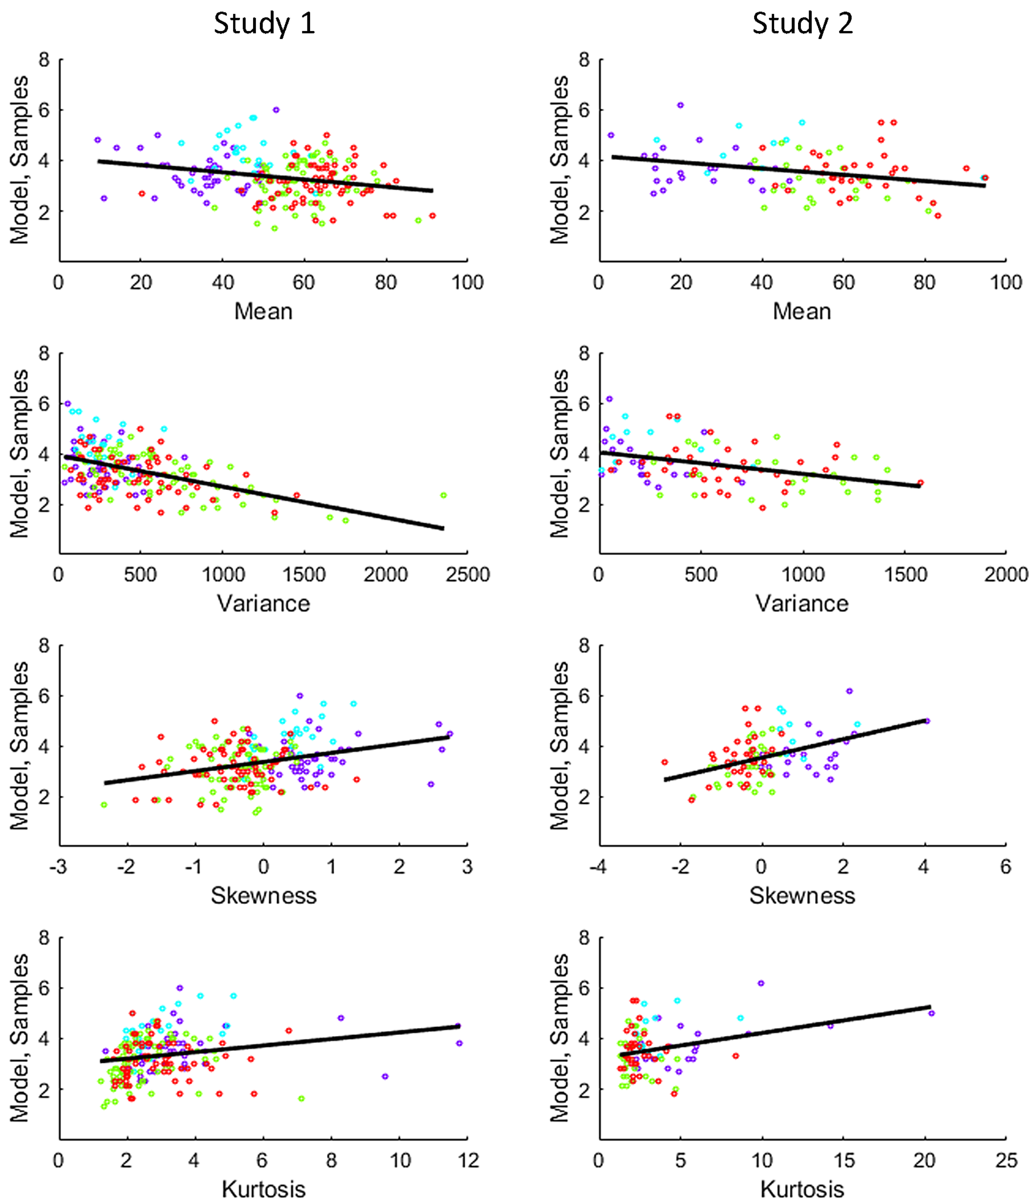

FIGURE 4.6: Scatterplots of generating distribution moments and the mean number of samples for the optimal model, for Study 1 and Study 2, separated by domain: male faces (purple), female faces (cyan), food (green), and holiday destinations (red). The black line is the regression line.

sampling rate. Our findings also go beyond those of previous studies which have mainly focused on altering the skewness of the generating distribution and neglected to investigate other distribution moments. In fact, upon closer investigation of the distributions used by Baumann et al. (2020) in Experiment 2 we found that the distributions did not just differ in skewness, but in mean, variance and kurtosis as well. As we found that participants' sampling rate could be predicted by both the mean and skewness of the generating distribution, we believe that no conclusive claims can be made solely regarding the relationship between the skewness of the generating distribution and sampling behaviour.

At this point, one might speculate that there might exist an unknown individual difference which could dispose some individuals who tend to evaluate options in a positively skewed way, to sample more as well. Although as yet there is no evidence for such a disposition, and we cannot know at this time what this disposition would be, the discovery of such a disposition would have important implications for predicting real-world decisions. That is, participants' decision patterns should be predictable from how they subjectively evaluate options. Although this explanation might be tempting, it is more likely that the link between the moments of the generating distribution and participants' sampling rate arises because of computational mechanisms involved inherently in solving optimal stopping problems. Indeed, we found that moment values produced similar sampling effects in our optimal model (Figure 4.6) as they did for the participants (Figure 4.5), even though the model has no individual disposition to sample more or less, and merely computes the solution to the optimal stopping problem. Moreover, we observe that the relationship between sampling rate and distribution moments does not hold for participants in general, but depends on which domain a participant was assigned to. For example, the participants who sampled the most were the ones in the positively skewed male face domain (see Figure 4.5).

Many previous studies have attempted to compute optimal measures based on an approximation of participants' perceived option generating distribution. Some studies (e.g., Cardinale et al., 2021; Costa & Averbeck, 2015) assumed participants used real-world values as their generating distribution and so offered them options sampled from real-world markets to conform with participants' presumed pre-existing prior. Other studies attempted to teach participants the generating distribution either through descriptions using statistical terminology and/or graphs of the probability densities of statistical distributions (Baumann et al., 2020; Lee & Courey, 2020), enriched feedback and/or financial rewards (Campbell & Lee, 2006), or through repeated interactions with the sequences of options (Goldstein et al., 2020). These kinds of learning schemes are most suited to scenarios where the rank of an option within its sequence can be computed directly from the objective (numerical) values.

However, for some real-world scenarios, many of these learning schemes might not be representative of real-life distribution learning, which occurs through repeated interaction with many options, often at random. Additionally, a participant's subjective valuation of a stimulus such as a number (e.g., a price), may not necessarily equal the number's objective value, as it is displayed on an exogenous statistical distribution provided by the researcher to a participant. For example, a participant might subjectively valuate a sizeable difference between £0 and £10 but a negligible difference between £1000 and £1010 and then act on this subjective valuation (rather than the objective price) when making decisions. This distinction would be relevant for full information problems, where the absolute option value (rather than its relative rank) is needed to solve the decision problem. Even further, the style with which participants valuate numbers or other stimuli is likely subject to individual differences. Methods that attempt to teach participants generating distributions using mathematical representations cannot account for such factors.

Therefore, an advantage of our study is that participants generated their own prior distribution of subjective values of attractiveness, which ensures that they know the generating distribution and that the optimality model operates on participants' personalised subjective values. Furthermore, our manipulation of decision-making domain effectively provided an experimental manipulation of distribution moments, but using moments that are representative of natural image-based domains. Although we did not vary the generating distribution's moments exogenously, our method is more applicable to real-world scenarios where option values, and consequently options' ranks, are assigned subjectively, e.g., in mate choice (Furl et al., 2019). Additionally, as more researchers may want to investigate sampling behaviour on optimal stopping tasks using images rather than numbers, our way of specifying the mean and variance of the generating distribution might yet be the best option. After all, using only images requires the researcher to obtain the option values separately for each individual, as the value of these kind of complex, naturalistic stimuli often cannot be objectively defined (Trendl et al., 2021).

To conclude, this paper provides novel insights into human sampling behaviour by directly comparing decision biases across three image-based decision making domains. Our studies support earlier findings on the facial attractiveness task, showing that participants oversample compared to an optimal model. Results for the decision-making domains food and holiday destinations also revealed an oversampling bias in participants. Additionally, and perhaps most importantly, we found evidence that sampling biases can be predicted by the mean and skewness of the underlying distribution of domain-specific option values.

References

Abdelaziz, F. B., & Krichen, S. (2006). Optimal stopping problems by two or more decision makers: A survey. *Computational Management Science*, *4*(2), 89–111. https://doi.org/10.1007/s10287-006-0029-5

Anwyl-Irvine, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2020). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods*, *52*(1), 388–407. https://doi.org/10.3758/s13428-019-01237-x

Backwell, P. R. Y., & Passmore, N. I. (1996). Time constraints and multiple choice criteria in the sampling behaviour and mate choice of the fiddler crab, Uca annulipes. *Behavioral Ecology and Sociobiology*, *38*(6), 407–416. https://doi.org/10.1007/s002650050258

Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23), 12750–12755. https://doi.org/10.1073/pnas.2002312117

Bearden, J. N., Rapoport, A., & Murphy, R. (2006). Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Science*, *52*(9), 1437–1449. https://doi.org/10.1287/mnsc.1060.0535

Blechert, J., Lender, A., Polk, S., Busch, N. A., & Ohla, K. (2019). Food-Pics_Extended — An Image Database for Experimental Research on Eating and Appetite: Additional Images, Normative Ratings and an Updated Review. *Frontiers in Psychology*, *10*, 307. https://doi.org/10.3389/fpsyg.2019.00307

Campbell, J., & Lee, M. D. (2006). The Effect of Feedback and Financial Reward on Human Performance Solving 'Secretary' Problems. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *28*(28).

Cardinale, E. M., Pagliaccio, D., Swetlitz, C., Grassie, H., Abend, R., Costa, V. D., Averbeck, B. B., Brotman, M. A., Pine, D. S., Leibenluft, E., & Kircanski, K. (2021). Deliberative Choice Strategies in Youths: Relevance to Transdiagnostic Anxiety Symptoms. *Clinical Psychological Science*, 1–11. https://doi.org/10.1177/2167702621991805

Costa, V. D, & Averbeck, B. B. (2015). Frontal-Parietal and Limbic-Striatal Activity Underlies Information Sampling in the Best Choice Problem. *Cerebral Cortex*, *25*(4), 972–982. https://doi.org/10.1093/cercor/bht286

Faul, F., Erdfelder, E., Buchner, A., & Lang, A.-G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, *41*(4), 1149–1160. https://doi.org/10.3758/BRM.41.4.1149

Ferguson, T. S. (1989). Who Solved the Secretary Problem? *Statistical Science*, *4*(3), 282–289. https://doi.org/10.1214/ss/1177012493

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

Furl, N., & Averbeck, B. B. (2011). Parietal Cortex and Insula Relate to Evidence Seeking Relevant to Reward-Related Decisions. *Journal of Neuroscience*, *31*(48), 17572–17582. https://doi.org/10.1523/JNEUROSCI.4236-11.2011

Furl, N., Averbeck, B. B., & McKay, R. T. (2019). Looking for Mr(s) Right: Decision bias can prevent us from finding the most attractive face. *Cognitive psychology*, *111*, 1–14. https://doi.org/10.1016/j.cogpsych.2019.02.002

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, *61*(313), 35–73. https://doi.org/10.2307/2283044

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning When to Stop Searching. *Management Science*, *66*(3), 1375–1394. https://doi.org/10.1287/mnsc.2018.3245

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339–361. https://doi.org/10.1037/dec0000081

Guan, M., Lee, M. D., & Silva, A. (2014). Threshold Models of Human Decision Making on Optimal Stopping Problems in Different Environments. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*(36).

Guan, M., & Stokes, R. (2020). A cognitive modeling analysis of risk in sequential choice tasks. *Judgment and Decision making*, *15*, 823–850.

Hauser, T. U., Moutoussis, M., Dayan, P., & Dolan, R. J. (2017). Increased decision thresholds trigger extended information gathering across the compulsivity spectrum. *Translational Psychiatry*, *7*(12), 1–10. https://doi.org/10.1038/s41398-017-0040-3

Hauser, T. U., Moutoussis, M., Purg, N., Dayan, P., & Dolan, R. J. (2018). Beta-Blocker Propranolol Modulates Decision Urgency During Sequential

Information Gathering. *The Journal of Neuroscience*, *38*(32), 7170–7178. https://doi.org/10.1523/JNEUROSCI.0192-18.2018

Hill, T. P. (2009). Knowing When to Stop: How to gamble if you must—the mathematics of optimal stopping. *American Scientist*, *97*(2), 126–133.

Ivy, T. M., & Sakaluk, S. K. (2007). Sequential mate choice in decorated crickets: Females use a fixed internal threshold in pre- and postcopulatory choice. *Animal Behaviour*, *74*(4), 1065–1072. https://doi.org/10.1016/j.anbehav.2007.01.017

Kahan, J. P., Rapoport, A., & Jones, L. V. (1967). Decision making in a sequential search task. *Perception & Psychophysics*, *2*(8), 374–376. https://doi.org/10.3758/BF03210074

Lee, M. D. (2006). A Hierarchical Bayesian Model of Human Decision-Making on an Optimal Stopping Problem. *Cognitive Science*, *30*(3), 1–26. https://doi.org/10.1207/s15516709cog0000_69

Lee, M. D., & Courey, K. A. (2020). Modeling Optimal Stopping in Changing Environments: A Case Study in Mate Selection. *Computational Brain & Behavior*, *4*(1), 1–17. https://doi.org/10.1007/s42113-020-00085-9

Matejka, J., Glueck, M., Grossman, T., & Fitzmaurice, G. (2016). The Effect of Visual Appearance on the Performance of Continuous Sliders and Visual Analogue Scales. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 5421–5432. https://doi.org/10.1145/2858036.2858063

MATLAB. (2015). Version 2015b. The MathWorks Inc. *Natick, Massachusetts*.

Milinski, M., & Bakker, T. (1992). Costs influences sequential mate choice in sticklebacks, gasterosteus aculeatus. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *250*(1329), 229–233. https://doi.org/10.1098/rspb.1992.0153

Morey, R. D., & Rouder, J. N. (2018). BayesFactor: Computation of Bayes Factors for Common Designs.

Prolific. (2014). Available at: https://www.prolific.co. *Oxford, UK*.

RStudioTeam. (2020). RStudio: Integrated Development Environment for R. *Boston, Massachusetts*.

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

Shapira, Z., & Venezia, I. (1981). Optional stopping on nonstationary series. *Organizational Behavior and Human Performance*, *27*(1), 32–49. https://doi.org/10.1016/0030-5073(81)90037-4

Shu, S. B. (2008). Future-biased search: The quest for the ideal. *Journal of Behavioral Decision Making*, *21*(4), 352–377. https://doi.org/10.1002/bdm.593

Taber, K. S. (2018). The Use of Cronbach's Alpha When Developing and Reporting Research Instruments in Science Education. *Research in Science Education*, *48*(6), 1273–1296. https://doi.org/10.1007/s11165-016-9602-2

Trendl, A., Stewart, N., & Mullett, T. L. (2021). A zero attraction effect in naturalistic choice. *Decision*, *8*(1), 55. https://doi.org/10.1037/dec0000145

Van der Leer, L., Hartig, B., Goldmanis, M., & McKay, R. (2015). Delusion proneness and 'jumping to conclusions': Relative and absolute effects. *Psychological Medicine*, *45*(6), 1253–1262. https://doi.org/10.1017/S0033291714002359

Wagenmakers, E.-J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Selker, R., Gronau, Q. F., Dropmann, D., Boutin, B., Meerhoff, F., Knight, P., Raj, A., van Kesteren, E.-J., van Doorn, J., Šmíra, M., Epskamp, S., Etz, A., Matzke, D., . . . Morey, R. D. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic Bulletin & Review*, *25*(1), 58–76. https://doi.org/10.3758/s13423-017-1323-7

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer Sequential Search: Not Enough or Too Much? *Marketing Science*, *22*(4), 503–519. https://doi.org/10.1287/mksc.22.4.503.24909

Supplementary materials

S1 Attention check Study 1

We decided to add an attention check to phase one of Study 1 to compensate for the unsupervised nature of online data collection. Each attention check comprised two screens that were shown one after the other. Attention checks showed up at nine random points (5% of the total of 180 images) throughout phase one. This totals 18 attention checks (nine time points x two screens). Each attention check screen showed a cross (either black or red), a 'next' button, and the text "press 'next' when the cross disappears". The cross disappeared at a random time interval between one and five seconds. The 'next' button was active the whole time. Reaction times for pressing the 'next' button were recorded for both screens, that is, for both the black and the red cross. Before data analysis, participants' response time for pressing the 'next' button was compared to the actual time interval before the cross disappeared (cross display time). If participants were paying attention, they would not press the 'next' button as soon as it appeared, but would instead read the text and respond only after the cross had disappeared. Thus, if participants' response time exceeded the cross display time, they passed the attention check.

S2 Filenames corresponding to food images

0001, 0002, 0004, 0007, 0009, 0010, 0016, 0022 ,0025, 0032, 0044, 0049, 0053, 0054, 0057, 0061, 0072, 0080, 0089, 0095, 0101, 0104, 0110, 0113, 0123, 0143, 0145, 0150, 0153, 0157, 0166, 0167, 0175, 0176, 0192, 0194, 0198, 0199, 0200, 0201, 0206, 0222, 0227, 0233, 0244, 0248, 0249, 0250, 0251, 0255, 0256, 0258, 0259, 0269, 0278, 0279, 0280, 0281, 0282, 0283, 0285, 0298, 0311, 0313, 0317, 0319, 0321, 0323, 0338, 0347, 0350, 0375, 0434, 0491, 0507, 0512, 0557, 0563, 0567, 0569, 0581, 0602, 0631, 0654, 0662, 0741, 0770, 0810, 0894, 0896

S3 Outlier removal Study 1

Although restrictions were set in Prolific to collect only 75 participants per domain, upon inspection of the data, we discovered that 76 participants were recruited in the food domain. As no duplicate IDs were found, we included all 76 participants in the data analysis. Also of note is that one participant's self-reported age was 16 (faces domain) and one participant's self-reported age was 36 (food domain), despite the enrolment restrictions set beforehand on Prolific. Considering that neither age required ethical reconsideration under British Psychological Society guidelines, we decided to include both participants in the analyses. To control for task incongruent behaviour, we pre-registered that all data points (i.e., mean number of samples and mean rank for each participant) had to be within 2.5 SD of each condition mean. We found three data points violating this assumption: one in the faces domain (in the rank of the chosen face), and two in the food domain (in the number of samples). These data points were thus excluded from the data analysis. If participants failed $>$ 25% of the attention checks (i.e., more than five) they were also excluded from the data analysis. Using this measure, another 18 participants were excluded (seven in the faces domain, four in the food domain, and seven in the holiday destinations domain).

S4 Outlier removal Study 2

One data point in the faces domain, for the mean number of samples, was excluded because it was $>$ 2.5 SD from the condition mean.

TABLE S1: Demographic statistics for each of the three domains: faces, food and holiday destinations. *One participant did not provide a valid response to this demographics question.

	Faces (N = 68)	**Food** (N = 72)	**Holiday Destinations** (N = 68)
Age			
Mean (SD)	26.43 (4.87)	25.53 (5.26)	27.13 (4.60)
Missing*	1	0	0
Sex			
Male	25	28	24
Female	42	43	43
Other	1	1	1
Nationality			
United Kingdom	34	44	54
Ireland	0	6	1
United States	14	11	7
Canada	19	11	3
Australia	0	0	3
New Zealand	1	0	0

TABLE S2: Demographic statistics for Study 2, for each of the three domains: faces, food and holiday destinations.

	Faces (N = 32)	**Food** (N = 28)	**Holiday Destinations** (N = 36)
Age			
Mean (SD)	28.13 (14.89)	35.68 (18.23)	28.22 (15.71)
Sex			
Male	5	11	10
Female	26	17	26
Other	1	0	0

Chapter 5

Methodological remarks regarding optimal stopping tasks and the implications for sampling biases

Abstract

This paper investigates a type of optimal stopping problem where options are presented in sequence and, once an option has been rejected, it is impossible to go back to it. With previous research finding mixed results of undersampling and oversampling biases on these kinds of optimal stopping tasks, the question remaining is what causes people to sample too much or too little compared to models of optimality? In two pilot studies and a main study, we explored task features that could lead to over- versus undersampling on number-based tasks. We found that, regardless of task features, there were no significant differences in human sampling rate across conditions. Nevertheless, we observed differences in sampling biases across conditions due to varying sampling rates of the optimal model. Our optimal model, like most models used for this type of optimal stopping problem, requires that researchers specify the mean and variance of a theoretical distribution, from which the options are generated. We show that different ways of specifying this generating distribution can lead to different model sampling rates, and consequently, differences in sampling biases. This highlights that a correct specification of the generating distribution is critical when investigating sampling biases on optimal stopping tasks.

1 Introduction

Oftentimes in everyday life, decisions have to be made regarding options presented in sequence, like when attempting to find the best deal on a certain product or service. When should someone stop evaluating new information and commit to a decision? This common real-life dilemma can be defined as an optimal stopping problem. There are many types of optimal stopping problems, but here we specifically look at full information best choice problems in which participants first learn the probability distribution that will generate their decision options (e.g., from experience in the real world or from within the paradigm itself). Then, option values from this generating distribution are presented in sequence (e.g., finding new deals on different websites), and a decision maker has to decide when to stop sampling and choose an option, under the condition that rejected options cannot be returned to later (e.g., because the deal has expired) (for a review, see Freeman, 1983). To do this successfully, the decision maker must balance the potential of improving on the current option against the risk of losing the best option if too many options are sampled (Furl et al., 2019).

Previous studies exploring decision-making on economic optimal stopping tasks have reported that decision makers primarily stop searching too early compared to models of optimality (*undersampling*) (Bearden et al., 2006; Cardinale et al., 2021; Costa & Averbeck, 2015; Guan et al., 2014; Seale & Rapoport, 1997; Sonnemans, 2000). However, there are also examples of specific optimal stopping tasks on which people sample too much (*oversampling*), such as when choosing a date (Furl et al., 2019). Despite these contradicting findings, relatively little progress has been made in terms of characterising under which circumstances humans undersample or oversample on optimal stopping tasks. The current paper addresses this question by investigating various methodological task features that may affect sampling biases in three separate studies. This is important in light of recent research suggesting that optimal stopping tasks might have a wider real-world application, for example as part of cognitive behavioural therapy in anxiety disorders (Cardinale et al., 2021), or even as a general measure of problem solving ability and psychometric intelligence (Lee et al., 2005). For these kinds of real-world applications to be realised, more uniform and standardised procedures for studying human behaviour on optimal stopping tasks are warranted.

Presently, numerous versions of optimal stopping tasks prevail, which complicate direct comparisons between studies. For instance, countless different stimuli are used across the literature to indicate the value of an option (e.g., Baumann et al., 2020; Cardinale et al., 2021; Costa & Averbeck, 2015; Furl et al., 2019; Goldstein et al., 2020; Guan & Stokes, 2020), but there is reason to suggest that the type of stimulus (e.g., numbers or images) might affect human sampling behaviour. Specifically, a study by Costa and Averbeck (2015) found that participants undersampled on an economic optimal stopping task compared to a Bayesian ideal observer model. This behaviour was reported for a selection of decision scenarios including buying a subway ticket, a television, and a diamond ring. Option values were presented numerically for all decision scenarios. On a similar optimal stopping task (the 'facial attractiveness task'), however, where option values could be derived from an image only, Furl et al. (2019) observed that participants oversampled compared to the Bayesian ideal observer model. The reason why only images were used on the so-called facial attractiveness paradigm employed by Furl et al. (2019) was because the task aimed to investigate mate choice decisions: participants were instructed to choose the most attractive face from a sequence of faces as their date. There are a number of task features on which these studies varied, but one of the key differences between the two paradigms is that Furl et al. (2019) used naturalistic image-based stimuli (images of faces) and Costa and Averbeck (2015) used more abstract, numerical stimuli (e.g., prices).

Therefore, the aim of our first two pilot studies was to determine whether numeric stimuli necessarily lead to undersampling. Pilot Study 1 aimed to replicate undersampling on a version of the economic optimal stopping task employed by Costa and Averbeck (2015), which used smartphone prices. Then, in Pilot Study 2, participants still sequentially encountered prices for a smartphone contract, but the task otherwise retained all the other task features of the facial attractiveness paradigm described by Furl et al. (2019). In other words, the only change compared to Furl et al. (2019) was that the images of faces were replaced with numerical smartphone prices. Because of the use of numbers instead of image-based stimuli, we hypothesised that this adaptation to the paradigm would be sufficient to induce an undersampling bias, in line with the results of previous studies that employed numerical stimuli (e.g., Baumann et al., 2020; Bearden & Connolly, 2007; Cardinale et al., 2021; Costa & Averbeck, 2015; Furl & Averbeck, 2011; Sonnemans, 2000).

However, Pilot Study 2 showed that participants oversampled on an economic number-based task that implemented task features of the facial attractiveness task (Furl et al., 2019). Therefore, the aim of our Main Study was to delineate which methodological task feature(s) could have led to oversampling on the number-based task. We hypothesised that at least one task feature implemented in Furl et al. (2019) and Pilot Study 2, which was not present in Costa and Averbeck (2015) and Pilot Study 1, might have been responsible for the observed oversampling bias. At this point, we were in a position to perform a Main Study that attempted to replicate Pilot Study 1 and Pilot Study 2, and added additional conditions to systematically isolate the task feature that leads to oversampling.

2 General Materials and Methods

Participants After excluding participants who did not pass the attention check (Supplementary Materials, text A), 390 participants were included across our three studies (N_{pilot1} = 50, N_{pilot2} = 46, N_{main} = 294). Participants were recruited through the online recruitment service Prolific (Prolific, 2014), and were all fluent in the English language. As our studies involved presenting participants with phone prices in GBP, we used Prolific's pre-screening facility to ensure that all participants were residents of the United Kingdom. Gorilla Experiment Builder (Anwyl-Irvine et al., 2020) was used to create and host the studies. Across all studies, participants were presented with an instruction screen prior to commencing the study, and informed consent was obtained in accordance with the Declaration of Helsinki. All three studies were approved by Royal Holloway, University of London's Ethics Board.

Stimuli Participants in all three studies were told that they were buying a new smartphone. They were presented with sequences of prices for flagship models by the top brands (e.g., iPhone, Samsung, Huawei), on an up to 5GB plan with unlimited texts and minutes. All prices were actual prices (in GBP) of 2-year contracts offered by various UK retailers as harvested from internet advertisements in the year before data collection. In this way, we attempted to approximate participants' real-world expectations of prices on the market as closely as possible.

Bayesian ideal observer model Human behaviour on our optimal stopping tasks was compared to a Bayesian ideal observer model, for which performance is Bayesian optimal. This computational Markov decision process (MDP) model has been used in previous literature, including Costa and Averbeck (2015), Furl et al. (2019) and Cardinale et al. (2021) (for a mathematical description of the model, see Supplementary Materials, Text B). Just like the historically used Gilbert and Mosteller model (Gilbert & Mosteller, 1966), the ideal observer model's expectations about future option values are based on a standard normal distribution, from which future options are assumed to be generated. Researchers using these types of model generally fix the mean and variance of this 'generating distribution' in advance to what they think participants are likely to use when making decisions. For the Bayesian ideal observer model, where the generating distribution is updated based on each new sample, researchers fix the mean and variance of the prior of the generating distribution (i.e., its initial value, before option sampling begins). Here, we set the prior of the generating distribution of the ideal observer model in two possible ways (Model 1 and Model 2), depending on the task features. These will be explained in more detail below.

Pilot Study 1 used the original MATLAB code (MATLAB, 2015) generously provided by Costa and Averbeck (2015). The version of the model we used in Pilot Study 1 (Model 1) received as input the same sequence values (i.e., phone prices) as the participants, in the order in which they were presented to the participants. Costa and Averbeck (2015), when implementing their ideal observer model, assumed that participants would use their experience with real world commodity prices when setting their prior distribution of option values. Costa and Averbeck (2015) therefore harvested commodity prices from real-world markets, and generated option sequences from these approximations to the real-world price distributions. We have done the same using smartphone prices that were also harvested from real-world markets. We are assuming that participants attempt to choose the option with the maximal subjective value, but that participants' subjective values of the options are equal to the options' exact (objective) price values which the model receives as input.

Like Costa and Averbeck (2015) and Cardinale et al. (2021), options were modeled as samples from a Gaussian distribution with a normal-inverse-χ^2 prior. The prior distribution has four parameters: the prior mean (μ_0), the degrees of freedom of the prior mean (κ_0), the prior variance (σ_0^2), and the degrees of freedom of the prior variance (ν_0). For each sequence, the values of μ_0 and σ_0^2 were set to the mean and variance of the log transformed distribution of raw phone prices (i.e., all 90 possible phone prices; μ_0 = -6.7402, σ_0^2 = 0.1038). Log transformation was applied to the prices to approximate normality: a Shapiro-Wilk test of normality indicated that phone prices were not normally distributed (W = 0.94, $p < .001$). Costa and Averbeck (2015) fixed the prior distribution in a slightly different way as we did, as they set the mean and variance of the model's prior generating distribution to that of each individual sequence's option values, rather than the whole distribution of option values. We tested whether this alternative specification of the prior of the generating distribution affected the model's sampling behaviour, but we found that the two similar ways of specifying the prior produced nearly identical sampling rates (Supplementary Materials, Figure S1). Model 1 employs a function R, which maps the rank of each option to its corresponding reward value. Reward values were assigned as follows: $R(1) = 0.12$, $R(2) = 0.08$, $R(3) = 0.04$, and $R(i > 3) = 0$, in accordance with the bonus payments that could be earned (see Section 3). As there was no explicit extrinsic cost-to-sample in the experimental design, the cost-to-sample parameter was fixed to zero.

Pilot Study 2 utilised a similar paradigm to Furl et al. (2019). Instead of assuming that participants use experience from the real world outside the study to set their prior, participants in Furl et al. (2019) learned the generating distribution within the study itself, and participants' subjective (reported) values of the stimuli were measured. Our participants were instructed to base their decisions on the optimal stopping task on their own distribution of attractiveness ratings (i.e., the subjective option values rather than the actual raw phone prices). This means that in the version of the model that we used for Pilot Study 2, the value of a given option in a sequence comprised the mean of participants' individual attractiveness ratings of that particular option in the rating phase. These mean ratings were put into the version of the model that we implemented for Pilot Study 2 (Model 2), in the same order in which they were presented to participants in the sequences. As outlined above, Costa and Averbeck (2015) modelled options as samples from a Gaussian distribution. To approximate a normal distribution in Pilot Study 2, ratings were log transformed for each participant before being put into the model: a Shapiro-Wilk test of normality indicated that attractiveness ratings were not normally distributed (W = 0.85, $p < .001$). In terms of the prior, Furl et al. (2019) set the mean and variance of the prior of their ideal observer model to those of the participants' subjective ratings of the stimuli in the generating distribution, which they learned prior to the optimal stopping task. We followed this procedure here by setting μ_0 and σ_0^2 to the mean and variance of the log transformed subjective value distribution (i.e., attractiveness ratings), which reflects the participant's and model's prior experience with the set of phone prices (Furl et al., 2019). The respective degrees of freedom for μ_0 and σ_0^2 were $\kappa_0 = 2$ and $\nu_0 = 1$. Reward values for Model 2 were set in the same way as Furl et al. (2019), meaning that we assumed that participants followed our instructions and tried to choose the option with the highest subjective value possible. Therefore, reward values were commensurate with the subjective value (attractiveness rating) of the chosen option. In other words, $R(1)$ = the subjective value of the highest ranked option, $R(2)$ = the subjective value of the second highest ranked item, and so on. The cost-to-sample parameter was fixed to zero because there was no explicit extrinsic cost-to-sample in the experimental design.

Conditions in our Main Study used either one of the two models outlined above, depending on the task design and instructions to participants (to be described in Section 5).

Data analysis The key dependent variable of interest for all three of our studies is the number of samples before choice (i.e., the position of the chosen price in the sequence). This variable is a mean value over the sequences for each participant.

The comparison of participants' sampling behaviour to the ideal observer model was done using MATLAB version 2015b (MATLAB, 2015) (repeated measures). Statistical tests were performed using RStudio (RStudioTeam, 2020). For all analyses, a p value of $< .05$ was considered significant.

3 Pilot Study 1

Experimental design Pilot Study 1 included 19 males, 30 females, and 1 participant who selected 'other' when reporting gender (M_{age} = 31.96, SD_{age}

= 10.67, range 18 to 65 years). Our design has been made openly available on Gorilla Open Materials [1]. Participants were presented with seven sequences of 12 prices each (Supplementary Materials, Figure S2). The order in which the sequences were presented was randomised in Gorilla. Costa and Averbeck (2015) rewarded participants financially for choosing one of the top three options in the sequence. In our study, participants were able to earn an additional £0.12 per sequence if they chose the lowest price, £0.08 if they chose the second lowest price, and £0.04 if they chose the third lowest price. Bonus payments were on top of a flat fee, which for all our studies was set in line with Prolific's recommended pay of at least £7.50 per hour. The paradigm utilised fixed screen timings, meaning that participants automatically advanced through the screens, except when asked to make a decision ('Take this option' or 'See next option'). Participants were warned about this feature in the instruction sheet.

Results and Discussion Recall from Section 2 that Model 1 uses a prior generating distribution with mean and variance calculated from the objective price distribution and attempts to maximise the monetary reward value of its choices. Contrary to our expectations, the comparison of participants' sampling rate to Model 1 did not replicate the undersampling bias reported by Costa and Averbeck (2015) and Cardinale et al. (2021). Instead, we found that there was no difference in sampling rate between participants and Model 1: $t(49) = -1.04$, $p = .302$ (Figure 5.1). The reason why Figure 5.1 shows no variation in mean values for Model 1 is because the order of the phone prices across the seven sequences was the same for each participant, and so the model always produced the same answer for these sequences. This characteristic means that the order of high quality and low quality options in a sequence could influence the mean sampling rate of Model 1 substantially, which might explain why we did not replicate undersampling. Because of these results, in our Main Study we employed multiple sequence orders to ensure we would obtain model results that are not specific to one particular sequence of options but rather an average over many sequences.

[1] https://gorilla.sc/openmaterials/53623

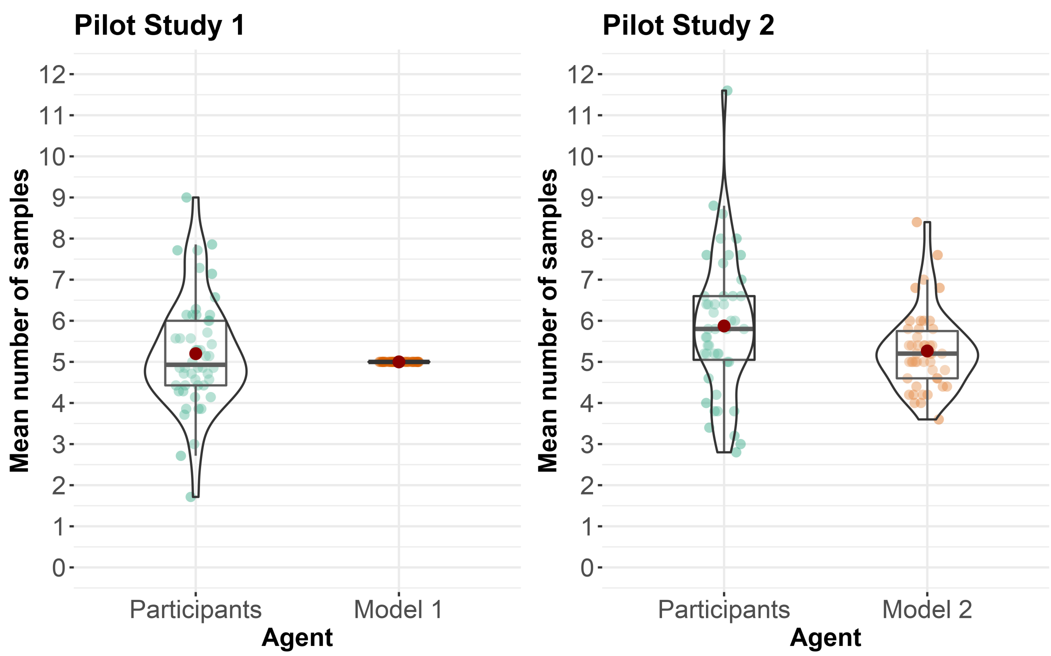

FIGURE 5.1: Distributions of the mean number of samples for participants versus Model 1 in Pilot Study 1, and participants versus Model 2 in Pilot Study 2. The red dots represent the mean, horizontal black lines represent the median, boxes show the 25% and 75% quantiles, and the whiskers represent the 95% confidence intervals.

4 Pilot Study 2

Experimental design We enrolled into Pilot Study 2 participants who did not participate in Pilot Study 1. Seventeen males and 29 females were included in our analysis of Pilot Study 2 (M_{age} = 30.57, SD_{age} = 11.36, range 18 to 75 years, four participants did not report their age). As with previous work (Furl et al., 2019), participants were presented with 180 prices (90 unique prices, all rated twice) in the first phase of the study. Prices were the same as used in Pilot Study 1. Phone prices appeared on the screen one at a time. Participants rated each price on its attractiveness using a slider scale from very unattractive (1) to very attractive (100). Attractiveness was defined as how willing participants were to buy this certain flagship model phone at the given price. Sliders on the slider scale were made invisible until first click to reduce slider biases (Matejka et al., 2016), and once clicked on, the slider showed the currently selected value on the scale. A progress bar was shown continuously at the bottom of the screen to visualise participants' progression.

Phase two of the study included five sequences. In each sequence, participants encountered 12 prices (Supplementary Materials, Figure S3). Smartphone prices were randomly sampled from the entire pool of prices that was rated in phase one. Participants were asked to attempt to choose as attractive a price as they could in every sequence, with the restriction that they could not return to a previously rejected price. The number of prices remaining in each sequence was shown at the top of the screen, and the rejected prices were shown at the bottom of the screen. When participants made a choice, they had to advance through a series of grey squares that replaced the remaining prices. This ensured that participants could not finish the study early by choosing an early option. Phase two was entirely self-paced - participants advanced by using their mouse to click on the buttons on the screen. If the last price in the sequence was reached, that price became their choice by default. After finishing a sequence, participants were directed to a feedback screen displaying their chosen price and the text: "This is the price of your contract! How rewarding is your choice?". Participants responded to this question using a slider scale ranging from not rewarding (1) to very rewarding (100). The feedback screen was included to provide feedback about the quality of the participants' choice by asking them to reflect upon its reward value before moving onto the next sequence, in lieu of the bonus payment screen in Pilot Study 1. Responses were not further analysed. Participants were reimbursed a flat fee only - no bonus payments were awarded.

Results and Discussion Recall from Section 2 that Model 2 uses a prior generating distribution calculated from the subjective values of the prices and attempts to maximise the subjective value of its choices. Because of the use of number-based stimuli, we hypothesised that participants would undersample compared to Model 2 in Pilot Study 2 where they searched for the most attractive smartphone price. However, we found that participants showed an oversampling bias instead: the comparison of participants' behaviour to the Model 2 version of the Bayesian ideal observer model showed that participants sampled significantly more options than Model 2 ($t(45) = 2.02$, $p < .05$; Figure 5.1). This result is in line with the results of Furl et al. (2019) on the facial attractiveness task, but contradicted our hypothesis. Although Pilot Study 1 and Pilot Study 2 used the same stimuli (smartphone prices), we found no evidence for sampling biases in Pilot Study 1, while participants showed an oversampling bias in Pilot Study 2. When directly comparing participants' sampling rates, we found that participants in Pilot

Study 2 sampled significantly more than participants in Pilot Study 1 ($t(86) = 2.14$, $p < .05$).

Hence, our results indicate that another task feature, rather than stimulus type, must account for the fact that we replicated oversampling in Pilot Study 2, despite not using images like previous research did (Furl et al., 2019).[2] Because statistical comparisons between studies where data were collected at different times should be treated with some caution, we will further investigate the difference in sampling rate between Pilot Study 1 and Pilot Study 2 by directly comparing these two paradigms in the same study (our Main Study).

We now highlight the key differences in task features between Pilot Study 1 and Pilot Study 2, which we further investigate in our Main Study. The first task feature that we will investigate is the rating phase that was included in Pilot Study 2. The aim of the rating phase was not only to obtain participants' subjective values for each of the prices, but also to familiarise them with the distribution of prices from which options in phase two are sampled. This could be crucial, as previous research has shown that participants are responsive to prior knowledge of varying generating distributions and adapt their sampling accordingly (Baumann et al., 2020; Guan & Lee, 2018; Guan et al., 2014). For example, Guan et al. (2014) reported that participants updated their decision thresholds in accordance with the quality of their environment (i.e., many high values/plentiful environment, many low values/scarce environment), while Baumann et al. (2020) found that participants sampled more in a scarce environment than in a plentiful environment. Although we attempted to match participants' expectations about how prices are distributed by including actual prices of UK retailers, participants in Pilot Study 1 could have been using different distributions based on their previous real-life experiences. As such, participants in Pilot Study 1 might have used different search strategies compared to participants in Pilot Study 2, who learned the underlying distribution we used for our study prior to commencing phase 2 of the task. Therefore, we consider the rating phase feature a strong contender in explaining participants' sampling biases.

A second possible influence of the rating phase is that subjective option values, rather than objective option values, can be used to determine the

[2]Oversampling biases on image-based optimal stopping tasks are also reported in Chapters 3 and 4.

highest ranking option in the sequence. The rating phase stems from Furl et al. (2019)'s facial attractiveness paradigm where it was essential to obtain each individual's personalised ratings for the faces that were presented in phase 2. Without the rating phase, the ranking of the faces could not have reflected each individual's true perception of facial attractiveness, as attractiveness is subjective. Thus, a certain face in theory could be the best option in a given sequence (rank 12) for one participant but the worst option (rank 1) for another participant. Keeping this in mind, it is possible that raw prices in Pilot Study 1 and subjective values in Pilot Study 2 were differently distributed because participants may not consider every GBP difference to be equal in subjective value. For example, a participant who believes any price of £800 or more is not worth choosing, might value a raw price of £800 and £900 in the same way, despite £800 being £100 cheaper and thus the better option. As such, the use of subjective values is another feature of the rating phase that makes the rating phase a contender for explaining participants' sampling biases.

There are additional differences in task features between Pilot Study 1 and Pilot Study 2, however, that must be considered. For example, after choosing an option, participants in Pilot Study 2 had to advance through a series of grey squares that replaced the remaining options. This feature was not incorporated in the previous implementations of the model that showed undersampling (Pilot Study 1; Cardinale et al., 2021; Costa & Averbeck, 2015). Although previous research has found no difference in sampling biases using versions with and without grey squares (Furl et al., 2019), the results have yet to be confirmed by directly contrasting a condition with grey squares with a matched condition without grey squares within the same study.

Furthermore, participants in Pilot Study 1 received bonus payments for choosing the lowest, second lowest, or third lowest price in the sequence, whereas participants in Pilot Study 2 were paid only the flat fee but were verbally instructed explicitly to maximise the subjective value of their choices. However, there is also evidence that awarding bonus payments for obtaining the best option in the sequence can actually increase sampling behaviour (Hsiao & Kemp, 2020). This seems inconsistent with the current results as we observe no increase in participants' sampling rate in Pilot Study 1 (which incorporated bonus payments) compared to Pilot Study 2 (no bonus payments). Therefore, further comparison between payoff structures is necessary to determine whether bonus payments might affect participants' sampling rate.

Finally, the pace of the two pilot studies was dissimilar, as Pilot Study 1 incorporated fixed timings for most of the screens, whilst Pilot Study 2 was entirely self-paced. The fixed timings in Pilot Study 1 effectively elongated the sequences, potentially giving participants a reason to choose sooner if they wanted to terminate the study earlier. This strategy would be less effective in a self-paced design like Pilot Study 2 where participants themselves decide how long they view an option. However, Pilot Study 2 was inherently a longer study than Pilot Study 1 due to the addition of the rating phase, which sheds doubt on the hypothesis that participants undersampled merely to end the study sooner. To determine whether the timing of the task could have influenced sampling biases, a direct comparison of an optimal stopping task with fixed timings and a self-paced optimal stopping task is warranted.

5 Main Study

Experimental design The differences between Pilot Study 1 and Pilot Study 2 (as outlined above) are further investigated in our Main Study, where we compare each of the task features directly to two control versions of the task, i.e., replications of Pilot Study 1 (*baseline* condition) and Pilot Study 2 (*full* condition). The other four conditions will henceforth be referred to as *squares*, *payoff*, *timing*, and *prior* (Table 5.1).

TABLE 5.1: Summary of condition characteristics for our Main Study.

		Condition					
		Baseline	Full	Squares	Payoff	Timing	Prior
Task feature	Grey squares		✓	✓			
	No bonus payments		✓		✓		
	Self-paced		✓			✓	
	Rating phase		✓				✓
Ideal observer	Model 1	✓		✓	✓	✓	✓
	Model 2		✓				

Demographic information for participants enrolled into each of our six conditions in our Main Study can be found in Table 5.2. Because of a technical difficulty with the participant recruitment platform, we overshot our data collection target in our Main Study by two participants, one in *timing* and one in *prior*. Participants across all conditions were presented with seven sequences of 12 prices each. Of note is that in Pilot Study 1, the order of the phone prices across the seven sequences was the same for each participant, which meant that there was no variation in the mean number of samples for the model (Figure 5.1). We were surprised to find a null result in Pilot Study 1 (see Figure 5.1) when we expected to replicate undersampling (Costa & Averbeck, 2015), and so we were concerned that the null result arose from the use of one stimulus sequence set that may or may not produce representative or generalisable behavioural performance. Therefore, in our Main Study, we strove to mitigate any such bias by introducing some variation in the model's performance. Hence, we created 10 different sets of seven sequences. Except for the *full* condition (i.e., the replication of Pilot Study 2), participants across all conditions were randomly assigned to one of the sets (fixed-ratio).

TABLE 5.2: Demographic statistics for each of the six conditions.

	Baseline (N = 50)	**Full** (N = 48)	**Squares** (N = 50)	**Payoff** (N = 51)	**Timing** (N = 50)	**Prior** (N = 45)
Age						
Mean (SD)	31.06 (10.63)	32.45 (12.58)	33.36 (10.40)	30.41 (11.82)	33.02 (11.66)	33.36 (12.39)
Missing data points	1	1	0	0	0	0
Sex						
Male	15	13	12	18	10	12
Female	34	33	38	32	39	33
Other	1	2	1	1	0	0
Prefer not to say	0	0	0	0	1	0

Baseline condition The first condition, henceforth referred to as *baseline*, was a redesigned version of Pilot Study 1 and attempted to replicate the undersampling bias reported on the economic optimal stopping task described in Costa and Averbeck (2015). Recall that participants were instructed to attempt to choose the lowest smartphone price in a sequence in order to maximise their earnings. The paradigm utilised fixed screen timings, and participants were able to earn bonus payments on top of the flat fee if they chose the lowest, second lowest, or third lowest price in the sequence (Supplementary Materials, Figure S2). As in Pilot Study 1, participants' sampling behaviour was compared to Model 1, which uses the full raw price distribution to set the mean and variance of the prior of the generating distribution.

Full condition The second condition attempted to replicate the oversampling bias observed in Pilot Study 2, and will henceforth be referred to as *full*. In this condition, participants first rated all possible phone prices on their attractiveness (phase 1), after which they commenced with the optimal stopping task (phase 2) where they were instructed to maximise the subjective value of their choices, that is, to choose the most attractive price in the sequence (Supplementary Materials, Figure S3). When participants made a choice, they had to advance through a series of grey squares that replaced the remaining prices. The entire paradigm was self-paced, and there were no bonus payments awarded on top of the flat fee. As in Pilot Study 2, participants' sampling behaviour was compared to Model 2, where the participants' subjective valuations of the prices are used to define the prior of the generating distribution and the option values.

Squares condition The third condition (*squares*) was the same as the *baseline* condition in that it was incentivised, had automatic timings, and did not use a rating phase. The only difference is that once participants had chosen an option in the *squares* condition (that was not the last option), they had to advance through the grey squares in a similar fashion to the *full* condition (Supplementary Materials, Figure S4), which was not the case in the *baseline* condition. Participants' sampling behaviour was compared to Model 1. If the task feature grey squares suffices to cause an oversampling bias, then we expect participants to sample more in the *squares* condition than in the *baseline* condition, leading to an oversampling bias in the *squares* condition but not in the *baseline* condition.

Payoff condition The fourth condition (*payoff*) was the same as the *baseline* condition in that it had no grey squares, had automatic timings, and did not use a rating phase. However, participants in the *payoff* condition did not receive any monetary bonus payments on top of the flat fee they received for their participation. Instead of receiving feedback regarding their earned bonus payments on the feedback screen, participants were shown pictures of either five stars, three stars or one star, if they chose respectively the lowest, second lowest, or third lowest price in the sequence (Supplementary Materials, Figure S5). Participants were specifically instructed that their goal was to maximise their number of stars. Therefore, reward values for Model 1 were changed to $R(1) = 5$, $R(2) = 3$, $R(3) = 1$, and $R(i > 3) = 0$, in line with the number of stars that participants could obtain. None of the other parameter values for Model 1 were changed. If the task feature no bonus payments suffices to cause an oversampling bias, then we expect participants to sample more in the *payoff* condition than in the *baseline* condition, leading to an oversampling bias in the *payoff* condition but not in the *baseline* condition.

Timing condition The fifth condition (*timing*) was the same as the *baseline* condition in that it had no grey squares, was incentivised, and did not use a rating phase. Instead of advancing through the screens of the optimal stopping task automatically, though, the *timing* condition incorporated a 'next' button in the top right corner of every option screen. This ensured that the entire paradigm was now self-paced. Participants' sampling behaviour was compared to Model 1. If the task feature self-paced suffices to cause an oversampling bias, then we expect participants to sample more in the *timing* condition than in the *baseline* condition, leading to an oversampling bias in the *timing* condition but not in the *baseline* condition.

Prior condition The sixth and final condition (*prior*) was the same as the *baseline* condition (no grey squares, incentivised, automatic timings) but added the rating phase of the *full* condition before the optimal stopping task. Although there was a phase 1 where participants expressed the subjective values of the distribution of potential options, the participants essentially ignored these phase 1 ratings in phase 2 and instead attempted to maximise their monetary bonus payment (i.e., by choosing the lowest phone price in the sequence which has the highest monetary payoff). As in the *baseline* condition, participants were able to earn bonus payments on top of the flat fee if they chose the lowest, second lowest, or third lowest price in the sequence. Participants' sampling behaviour was compared to Model 1 because participants attempted to maximise the monetary reward of their choices and not the subjective values from phase 1. If the task feature rating phase suffices to cause an oversampling bias, then we expect participants to sample more in the *prior* condition than in the *baseline* condition, leading to an oversampling bias in the *prior* condition but not in the *baseline* condition.

Results A 6x2 factorial ANOVA was used to compare the differential effects of our two agents (participants and model) across the six conditions. This analysis showed that there was a significant main effect of condition ($F(5,576) = 3.39$, $p < .01$), as well as a significant main effect of agent ($F(2,576)$

= 39.73, $p < .001$), as can be observed in Figure 5.2. However, despite the apparent differences in sampling bias between conditions (see Figure 5.2), we did not find a significant interaction effect of agent*condition ($F(4,576) = .90$, $p = .463$). Following this result, we wanted to assess whether the condition affected the participants' mean number of samples, as appeared to be the case for Pilot Studies 1 and 2. Human participant data (excluding the models) was analysed using Tukey's Honest Significant Difference (HSD) method. The results are shown in Table 5.3, and indicate that there was no evidence that participants sampled more options in any condition than any other. Therefore, when participants' sampling was directly contrasted within one study, the significant difference in sampling that arose between Pilot Study 1 and Pilot Study 2 did not replicate.

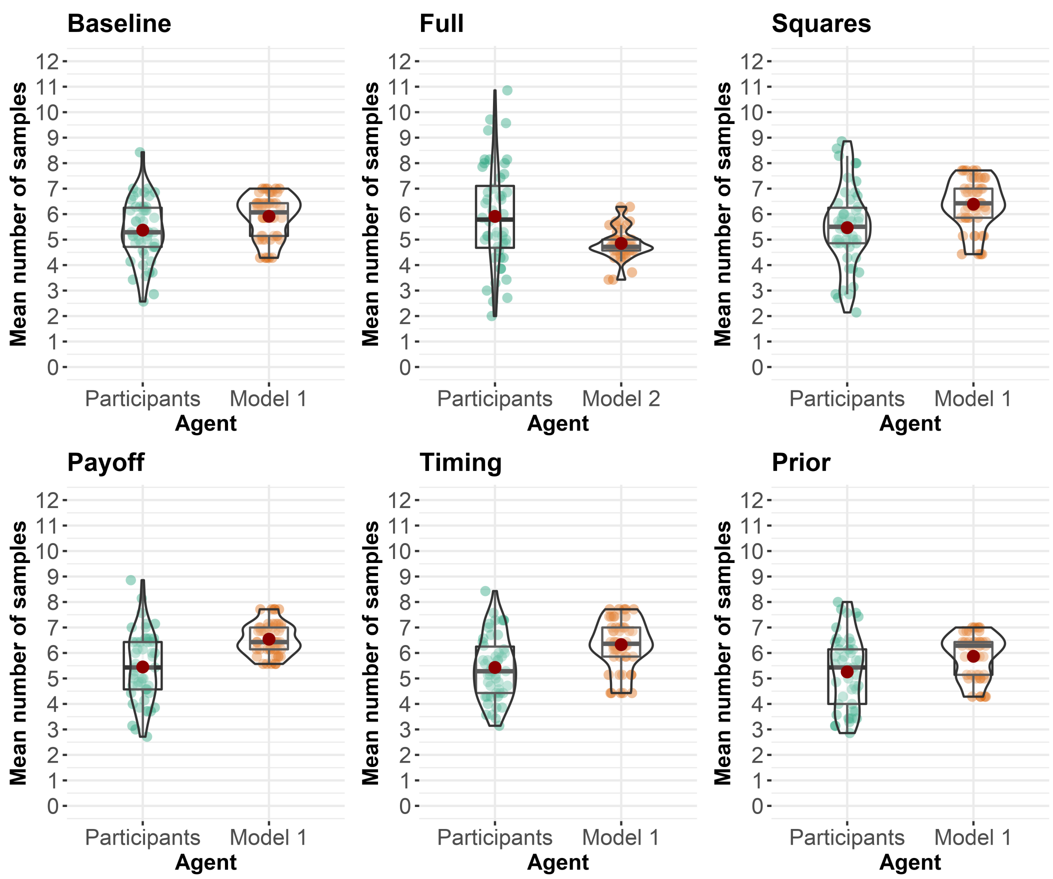

FIGURE 5.2: Distributions of the mean number of samples for participants versus their corresponding models, grouped by condition. The red dots represent the mean, horizontal black lines represent the median, boxes show the 25% and 75% quantiles, and the whiskers represent the 95% confidence intervals.

To test for differences in the mean number of samples between participants and the model, we performed post hoc pairwise t-tests (Bonferroni

TABLE 5.3: Adjusted p values indicating differences between the mean number of samples for participants across the six conditions. $p < .05$ is considered significant.

	Baseline	Full	Squares	Payoff	Timing	Prior
Baseline						
Full	.73					
Squares	~ 1	.93				
Payoff	~ 1	.91	~ 1			
Timing	~ 1	.87	~ 1	~ 1		
Prior	~ 1	.43	~ 1	~ 1	~ 1	

corrected for the six conditions) for each of the six conditions separately. Recall that in the *baseline*, *squares*, *payoff*, *timing* and *prior* conditions, the mean and variance of the prior of the generating distribution are set to those of the full distribution of raw phone prices (Model 1; Table 5.1), whereas for the *full* condition, the mean and variance of the prior of the generating distribution are set to those of the distribution of subjective values (Model 2; Table 5.1). The results of our post hoc analysis showed that in conditions using Model 1 (i.e., *baseline*, *squares*, *payoff*, *timing* and *prior*), participants undersampled ($t(49)$ = -2.57, $p < .05$, $t(49)$ = -3.36, $p < .01$, $t(50)$ = -4.84, $p < .001$, $t(49)$ = -4.25, $p < .001$, and $t(44)$ = -2.50, $p < .05$, respectively; Figure 5.2). In the *full* condition, which used Model 2, participants oversampled ($t(47)$ = 3.32, $p <$ .01; Figure 5.2).

Discussion In our Main Study, we investigated whether four candidate task features lead to oversampling on an economic optimal stopping task. The task features examined were grey squares (*squares*), no bonus payments (*payoff*), self-paced (*timing*) and rating phase (*prior*). Also included in our Main Study were a *baseline* condition, a redesigned version of Pilot Study 1, and a *full* condition, which attempted to replicate Pilot Study 2. Our results showed that participants undersampled in the *baseline*, *squares*, *payoff*, *timing* and *prior* conditions. This indicates that adding grey squares to the sequences, just paying participants a flat fee, having a self-paced task design, or adding a rating phase, does not affect human sampling biases on optimal stopping tasks. This was in contrast with our expectations, as we hypothesised that at least one of the candidate task features would lead to oversampling. We did replicate the oversampling bias of Pilot Study 2 in the *full* condition, bolstering our finding that the type of stimulus (numbers or images) alone cannot account for different sampling biases. We will now discuss an alternative theory to explain our findings.

Initially, we hypothesised that specific task features, and particularly the rating phase, might affect how humans sample on an optimal stopping task. Surprisingly, even though participants in our Main Study were presented with a diversity of task features across very different paradigms, we found no significant differences in human sampling rates across the six conditions. Instead, what caused sampling biases to differ was the behaviour of the Bayesian ideal observer model. Specifically, the model changed its optimal strategy depending on whether the prior of its generating distribution was set using the moments taken from the objective value distribution (raw prices) or the subjective value distribution (ratings). This highlights that if participants' generating distribution is unknown or incorrectly specified, apparent sampling biases could arise not because participants behave differently, but because the generating distribution the model operates on might be erroneous. We demonstrate this in Figure S6 in the Supplementary Materials: comparing participants in Pilot Study 2 and the *full* condition to Model 1 rather than Model 2 appears to flip our original results, causing a (slight) undersampling bias instead. Moreover, comparing participants in the *prior* condition to Model 2 rather than Model 1 resulted in no sampling bias, rather than the originally reported undersampling bias. This illustrates the need for standardised procedures for studying human behaviour on optimal stopping problems when using models that operate on a generating distribution (like the Gilbert and Mosteller model and the Bayesian ideal observer model). For example, one might wish to manipulate or control the (otherwise unknown) generating distribution so it can be modelled properly.

Previous research has tried different approaches to specify the prior participants operate upon in optimal stopping tasks. Baumann et al. (2020), for example, included a learning phase prior to the optimal stopping task to ensure that participants were acquainted with the generating distribution. Their learning phase encompassed the visual presentation of abstract mathematical representations of probability distributions. At the end, participants were asked to draw a histogram on which they received feedback. According to Goldstein and Rothschild (2014), such a graphical elicitation technique can lead to rather accurate representations of probability distributions in participants. Nevertheless, it is unlikely that people learn statistical distributions of options in the real world (e.g., when renting an apartment, or buying a smartphone) by memorising images of statistical distributions. Instead, they are more likely to build up a distribution from frequent sequential encounters. The assumption that this kind of learning happens in the real world formed the basis for the optimal model used in Costa and Averbeck (2015) and Cardinale et al. (2021), and our Model 1 as applied in Pilot Study 1 and the *baseline*, *squares*, *payoff* and *timing* conditions. In our *prior* condition, we provided a type of simulation of real-world sequential encounters with option values through the addition of a rating phase, thus ensuring that participants had learned the generating distribution of raw prices (which was otherwise implicit) prior to phase 2. Another study that incorporated learning is Goldstein et al. (2017), where participants learned an unknown distribution through repeated play. For optimal stopping tasks where the generating distribution is known to the researcher but unknown to the participants (e.g., as in number-based optimal stopping tasks like Pilot Study 1), any of the approaches discussed above might be used. Future research may wish to investigate which approach leads to the most accurate specification of participants' prior distribution, and thereby advise on a standardised procedure. For situations where the generating distribution is unknown to both the researcher and the participants (e.g., all image-based optimal stopping tasks), a rating phase which captures participants' subjective values, as incorporated in our Pilot Study 2, our *full* condition and Furl et al. (2019), might provide a solution. The main advantage of using subjective values is that the models' generating distribution can be unique for each participant. Participants can have different subjective values about options, and in this way, the model would be sensitive to these variations also.

Despite individual differences in subjective values, options' relative ranks should largely be preserved when using subjective values to set the mean and variance of the generating distribution. In our scenario of smartphone prices, the lowest price is also likely to be the highest rated price, thus both schemes should result in the same best-ranked item. This intuition was confirmed when we mapped the subjective attractiveness values as rated by participants in the *prior* condition onto the actual raw prices (Figure S7), which showed that the lowest smartphone prices received the highest subjective values. However, using subjective values instead of objective values is likely to affect the spacing between options, that is, two options with two different objective values might be viewed as similarly attractive by a participant. This is illustrated by the nonlinear relationship between objective and subjective values in Figure S7: participants make relatively small distinctions (the function appears flat) between objectively the lowest and highest prices, and participants' subjective evaluations primarily discriminate among intermediate prices. When thinking about real-life decision-making scenarios, this seems like an accurate representation of human decision-making: rarely will someone pass on a current smartphone deal if they subjectively perceive a potential future deal to be only incrementally better. Notably, this kind of subjective evaluation of option values could affect the shape of the generating distribution as well, which is known to have an influence on participants' sampling rate (Baumann et al., 2020; Guan & Lee, 2018; Guan et al., 2014). Figure S8 in the Supplementary Materials shows density plots of all participants' subjective attractiveness ratings recorded for this paper, i.e., in Pilot Study 1, the *full* condition and the *prior* condition, as well as a density plot of the full distribution of raw prices (objective values). Upon visual inspection, we can confirm that the distribution of objective values differs in shape from the distributions of subjective values, which could explain the reported differences in sampling biases between conditions.

Building on the fact that we found a nonlinear relationship between objective and subjective values, an interesting next step for future research could be to directly investigate how these kind of scaling effects influence model performance, beyond the inferences made from the existing dataset. Presently, our findings regarding the differences in distribution shapes are suggestive that the distance between option values could have an effect on the model. That is, as mentioned above, the distance between participants' subjective evaluations is greater for intermediate prices compared to the lowest and highest prices. We note here that it would be possible to look either at the additive effect, as we have done in this paper, or the proportional effect (i.e., the percentage change rather than the raw numbers). This distinction does not directly address the main hypothesis of the current paper as oversampling results are regardless of being additive or proportional. Nevertheless, future research may wish to look into this distinction to determine whether there are any differences when looking at scaling effects. This could be done, for example, in the context of a simulation study. Through model simulations, the scaling between subjective and objective values can be made explicit which allows for a high-quality investigation of scaling effects, whilst empirical research like the present study is somewhat limited in this regard.

One possible limitation is that besides the differences in the specification of the generating distribution between Model 1 and Model 2, the two models also incorporated a slightly different payoff structure, in line with the task design and instructions given to participants. We investigate the effect of varying the reward function on the sampling rate of Model 1 and Model 2 in the Supplementary Materials (Figure S9, Text C). Our supplementary results confirm that the difference in sampling of the models can best be explained by the different specification of the generating distribution.

6 Conclusion

Through three separate studies, we were able to show that none of the following task features significantly influenced participants' sampling rate on an optimal stopping task: use of images, adding grey squares, removing bonus payments, making the task self-paced, and adding a rating phase. In other words, these features cannot explain participants' sampling biases on optimal stopping tasks. Instead, we suggest that a correct specification of the generating distribution of option values is critical when investigating sampling biases on optimal stopping tasks, and several approaches to this challenge are discussed.

References

Anwyl-Irvine, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2020). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods*, *52*(1), 388–407. https://doi.org/10.3758/s13428-019-01237-x

Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23), 12750–12755. https://doi.org/10.1073/pnas.2002312117

Bearden, J. N., & Connolly, T. (2007). Multi-attribute sequential search. *Organizational Behavior and Human Decision Processes*, *103*(1), 147–158. https://doi.org/10.1016/j.obhdp.2006.10.006

Bearden, J. N., Rapoport, A., & Murphy, R. (2006). Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Science*, *52*(9), 1437–1449. https://doi.org/10.1287/mnsc.1060.0535

Cardinale, E. M., Pagliaccio, D., Swetlitz, C., Grassie, H., Abend, R., Costa, V. D., Averbeck, B. B., Brotman, M. A., Pine, D. S., Leibenluft, E., & Kircanski, K. (2021). Deliberative Choice Strategies in Youths: Relevance to Transdiagnostic Anxiety Symptoms. *Clinical Psychological Science*, 1–11. https://doi.org/10.1177/2167702621991805

Costa, V. D, & Averbeck, B. B. (2015). Frontal-Parietal and Limbic-Striatal Activity Underlies Information Sampling in the Best Choice Problem. *Cerebral Cortex*, *25*(4), 972–982. https://doi.org/10.1093/cercor/bht286

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

Furl, N., & Averbeck, B. B. (2011). Parietal Cortex and Insula Relate to Evidence Seeking Relevant to Reward-Related Decisions. *Journal of Neuroscience*, *31*(48), 17572–17582. https://doi.org/10.1523/JNEUROSCI.4236-11.2011

Furl, N., Averbeck, B. B., & McKay, R. T. (2019). Looking for Mr(s) Right: Decision bias can prevent us from finding the most attractive face. *Cognitive psychology*, *111*, 1–14. https://doi.org/10.1016/j.cogpsych.2019.02.002

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, *61*(313), 35–73. https://doi.org/10.2307/2283044

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2017). Learning in the Repeated Secretary Problem. *arXiv preprint arXiv:1708.08831*, arxiv 1708.08831.

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning When to Stop Searching. *Management Science*, *66*(3), 1375–1394. https://doi.org/10.1287/mnsc.2018.3245

Goldstein, D. G., & Rothschild, D. (2014). Lay understanding of probability distributions. *Judgment and Decision Making*, *9*(1), 1–14.

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339–361. https://doi.org/10.1037/dec0000081

Guan, M., Lee, M. D., & Silva, A. (2014). Threshold Models of Human Decision Making on Optimal Stopping Problems in Different Environments. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*(36).

Guan, M., & Stokes, R. (2020). A cognitive modeling analysis of risk in sequential choice tasks. *Judgment and Decision making*, *15*, 823–850.

Hsiao, Y.-C., & Kemp, S. (2020). The effect of incentive structure on search in the secretary problem. *Judgment and Decision Making*, *15*(1), 82–92.

Lee, M. D., O'Connor, T. A., & Welsh, M. B. (2005). Decision-Making on the Full Information Secretary Problem. *Proceedings of the Twenty-Sixth Conference of the Cognitive Science Society*, 819–824.

Matejka, J., Glueck, M., Grossman, T., & Fitzmaurice, G. (2016). The Effect of Visual Appearance on the Performance of Continuous Sliders and Visual Analogue Scales. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 5421–5432. https://doi.org/10.1145/2858036.2858063

MATLAB. (2015). Version 2015b. The MathWorks Inc. *Natick, Massachusetts*.

Prolific. (2014). Available at: https://www.prolific.co. *Oxford, UK*.

RStudioTeam. (2020). RStudio: Integrated Development Environment for R. *Boston, Massachusetts*.

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

Sonnemans, J. (2000). Decisions and strategies in a sequential search experiment. *Journal of Economic Psychology*, *21*(1), 91–102. https://doi.org/10.1016/S0167-4870(99)00038-0

Supplementary materials

Text A: Attention check

Multiple attention checks were added to phase one of Pilot Study 2 and the *full* and *prior* conditions in our Main Study, to compensate for the unsupervised nature of online data collection. Every attention check showed a cross, a 'next' button, and the text "press 'next' when the cross disappears". The cross disappeared at a random time interval between one and five seconds. The 'next' button was active the whole time. If participants were paying attention, they would not press the 'next' button as soon as it appeared, but would instead read the text and respond only after the cross had disappeared. Thus, if participants' response time exceeded the cross display time, they passed the attention check. Participants who failed $> 25\%$ of the attention checks were excluded from analysis. Four participants failed the attention check in Pilot Study 2, two failed in the *full* condition of our Main Study, and six failed in the *prior* condition of our Main Study. These participants were subsequently excluded from analysis.

Text B: Mathematical description of the Bayesian ideal observer model

Participants' sampling behaviour was compared to a Bayesian ideal observer model, where performance is Bayesian optimal and the cost-to-sample parameter was fixed to zero. Mathematically, the model is based on a discrete time Markov decision process (MDP) with continuous states (Cardinale et al., 2021; Costa & Averbeck, 2015). The MDP framework models the utility u of the state s at sample t as

$$u_t\left(s_t\right)=\max_{a \in A_{s_t}}\left\{r_t\left(s_t, a\right)+\int_s p_t\left(j \mid s_t, a\right) u_{t+1}(j) d j\right\} \tag{5.1}$$

where A_{s_t} is the set of available actions in state s at sample t. The term inside the curly brackets is the action value, and $r_t\left(s_t, a\right)$ is the reward that will be obtained in state s at sample t if action a is taken. The integral is taken over the set of possible states subsequent to the current sample. This set is weighted by the probability of transitioning into each state from the current state, given by $p_t\left(j \mid s_t, a\right)$. In other words, the utility u of the state s at sample t is the value of the best action a, which depends on reward value r, the cost-to-sample C_s, and the probabilities of outcomes j of subsequent states, weighted by their utilities.

At each position in the sequence, the ideal observer model computes the respective values for choosing the option and declining the option, and chooses the one with the highest value. To calculate the value of either taking or declining an option in a sequence, the model computes the action value Q as:

$$\begin{aligned} & Q_t\left(s_t, a=take\right)=r_t\left(s_t, a\right) \\ & Q_t\left(s_t, a=decline\right)=\int_s p_t\left(j \mid s_t, a\right) u_{t+1}(j) d j \end{aligned} \tag{5.2}$$

The key computations for the ideal observer model, as seen in equation 5.2, are utility (equation 5.1) and reward values (equation 5.3). The model uses backwards induction to derive utilities that could result from further sampling (equation 5.4).

$$
\begin{aligned}
r_t\left(s_t, a = accept\right) &= \sum_{i=1}^{N} p\left(rank = i\right) * R\left(i + \left(h - 1\right)\right) \\
r_t\left(s_t, a = decline\right) &= C_s
\end{aligned}
\tag{5.3}
$$

Function R assigns reward values to each option based on their rank, as further specified in Section 2 of the main manuscript for the Model 1 and Model 2 versions of the Bayesian ideal observer model. h represents the relative rank of the current option. When considering final sequence position N, the model computes final utilities as:

$$
u_N\left(s_N\right) = r\left(s_N\right) \text{ for all } s_N \in N \tag{5.4}
$$

and working backwards from N, we use equation 5.1 to compute utilities at every sequence position t.

Text C: Supplementary discussion of Figure S9

Visual comparison between the two reward functions plotted in Figure S9 suggests that Reward 1 increases the model's sampling rate, while Reward 2 decreases the model's sampling rate, for both Prior 1 and Prior 2. This finding further supports the need for standardised procedures, as sampling biases reported on optimal stopping tasks with different reward structures may not be directly comparable. Yet the reward function of the model alone cannot fully explain the observed differences in sampling biases between Model 1 and Model 2, as seen in Figure S6. For example, Prior 1 with Reward 2 has a fairly similar sampling rate compared to participants in both conditions, but Prior 2 with Reward 2 results in a clear oversampling bias. This difference in sampling of the models can only be explained by the different specification of the generating distribution. As such, it is the Prior 1 versus Prior 2 nature of Model 1 and Model 2 that accounts for their performance, rather than Reward 1 versus Reward 2.

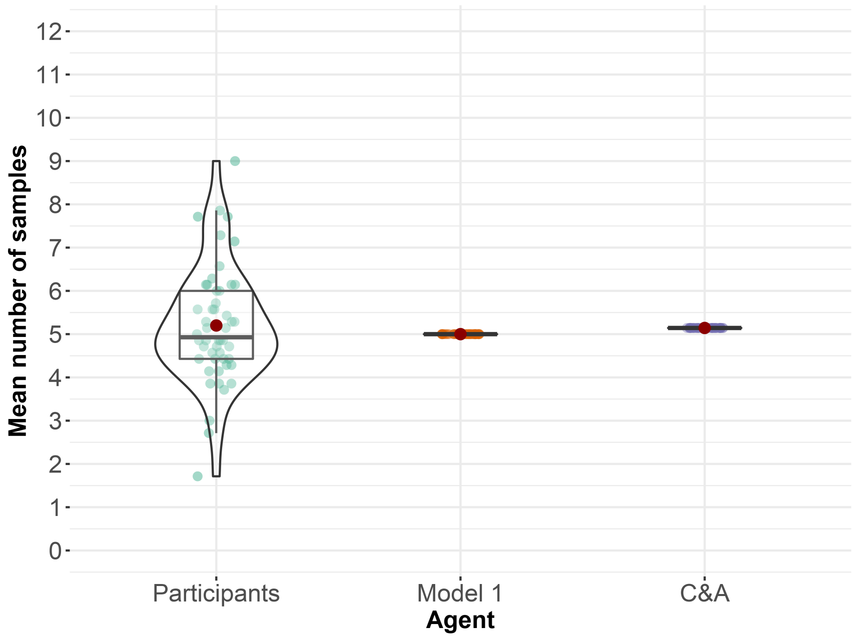

FIGURE S1: Comparison between the mean number of samples for our Model 1 and Costa & Averbeck's version of this model (C&A), using our results for Pilot Study 1. For the C&A model, the mean and variance of the model's prior generating distribution are set to that of each individual sequence's option values, rather than the whole distribution of option values (as in Model 1). The red dots represent the mean, horizontal black lines represent the median, boxes show the 25% and 75% quantiles, and the whiskers represent the 95% confidence intervals.

FIGURE S2: Screenshots of Pilot Study 1. A: instructions screen, shown for six seconds. B: option screen, shown for four seconds. C: choice screen. D: feedback screen where the participant chose the lowest price in the sequence, shown for three seconds.

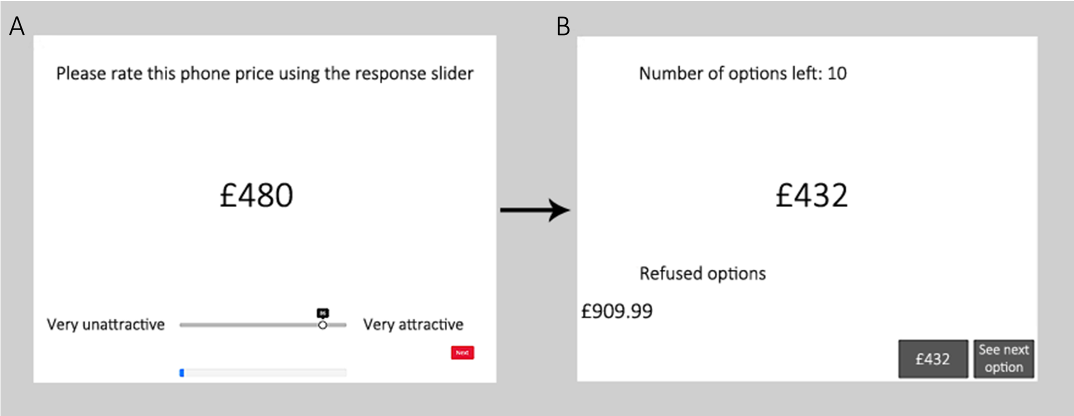

FIGURE S3: Screenshots of Pilot Study 2. A: rating screen. B: option screen.

FIGURE S4: Screenshots of our Main Study, condition *squares*. A: instructions screen, shown for six seconds. B: option screen, shown for four seconds. C: choice screen. D: grey square. E: feedback screen where the participant chose the lowest price in the sequence, shown for three seconds.

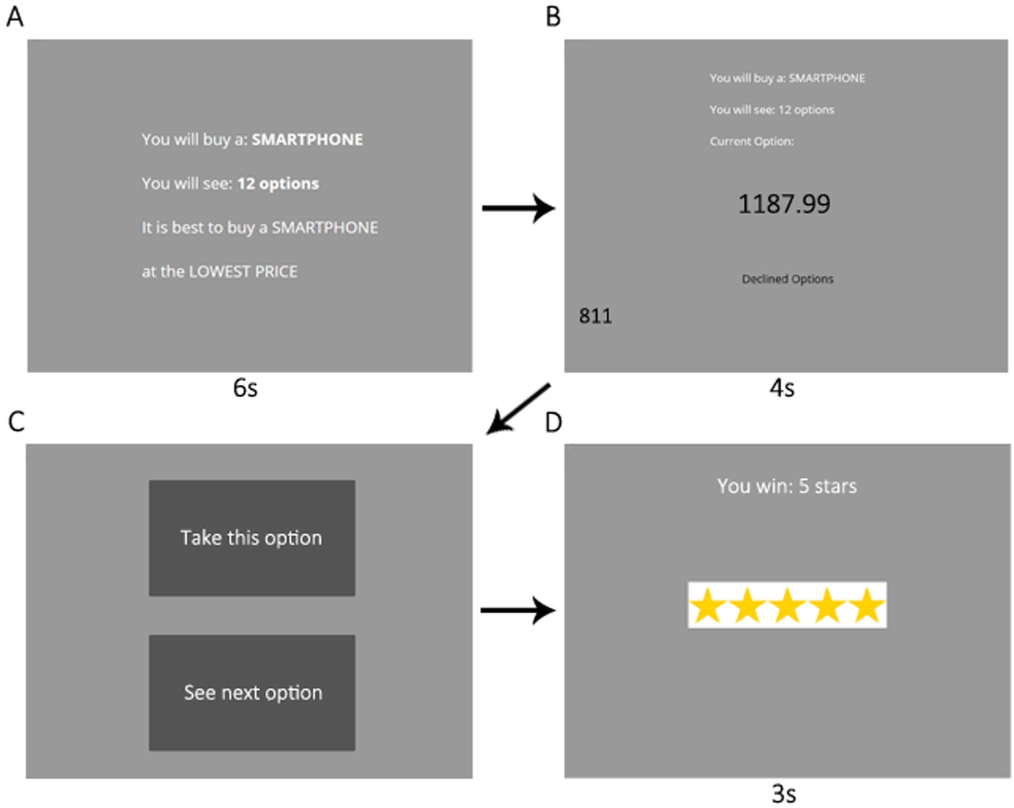

FIGURE S5: Screenshots of our Main Study, condition *payoff*. A: instructions screen, shown for six seconds. B: option screen, shown for four seconds. C: choice screen. D: feedback screen where the participant chose the lowest price in the sequence, shown for three seconds.

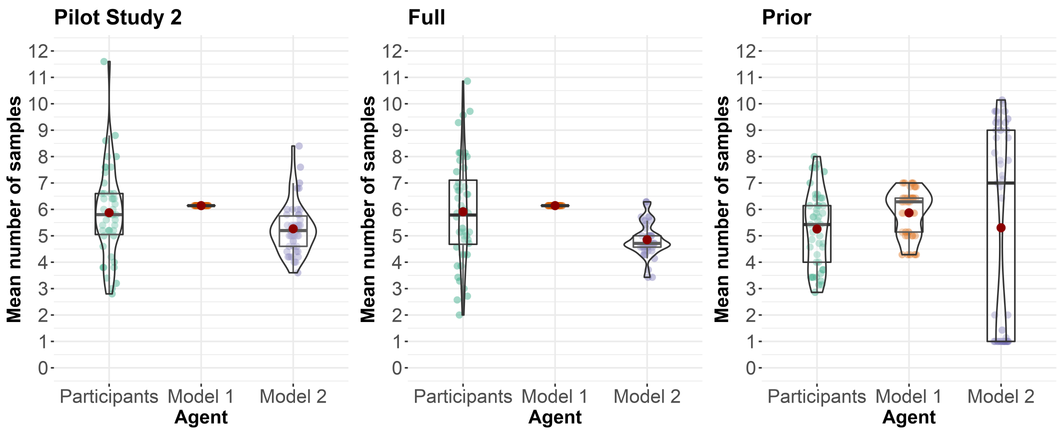

FIGURE S6: Distributions of the mean number of samples for participants versus Model 1 and Model 2 for Pilot Study 2, the *full* condition and the *prior* condition. The current comparison plot highlights how a different specification of the generation distribution can lead to different sampling biases within the same task. Specifically, for all three datasets, Model 1 tends to show at least a slight undersampling bias while Model 2 tends to show at least a slight oversampling bias or no sampling bias. There is no variation in Model 1's performance in Pilot Study 2 and the *full* condition because these conditions used only one set of sequences, rather than 10 like the *prior* condition. The red dots represent the mean, horizontal black lines represent the median, boxes show the 25% and 75% quantiles, and the whiskers represent the 95% confidence intervals. Reward functions of Model 1 and Model 2 were adapted slightly to reflect the instructions to participants in the relevant condition.

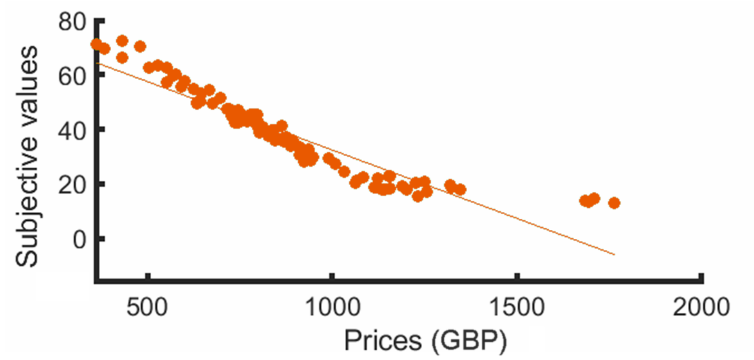

FIGURE S7: Nonlinear relationship between subjective attractiveness values and objective raw prices (GBP). Orange dots represent the mean subjective value for each price, the orange line represents the regression line. Data were taken from the *prior* condition.

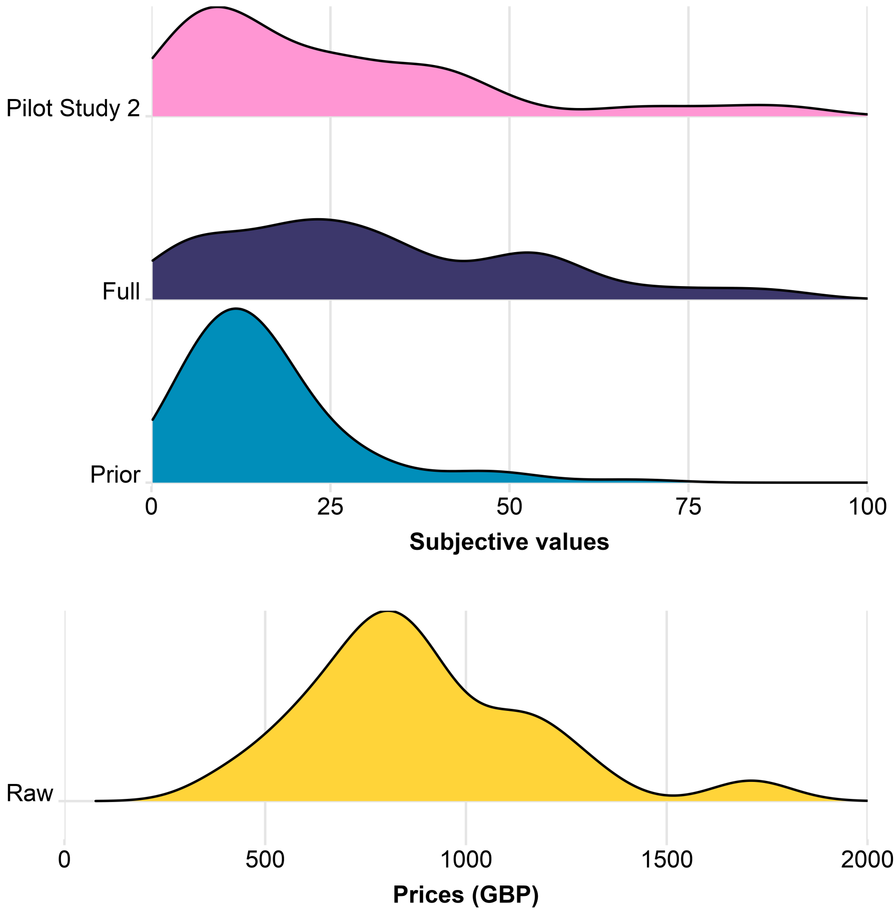

FIGURE S8: Top: density plots of participants' subjective values as collected in Pilot Study 2, the *full* condition and the *prior* condition. Bottom: density plot of the full distribution of raw prices (GBP) as used in Pilot Study 1, and the *baseline*, *squares*, *payoff*, *timing* and *prior* conditions.

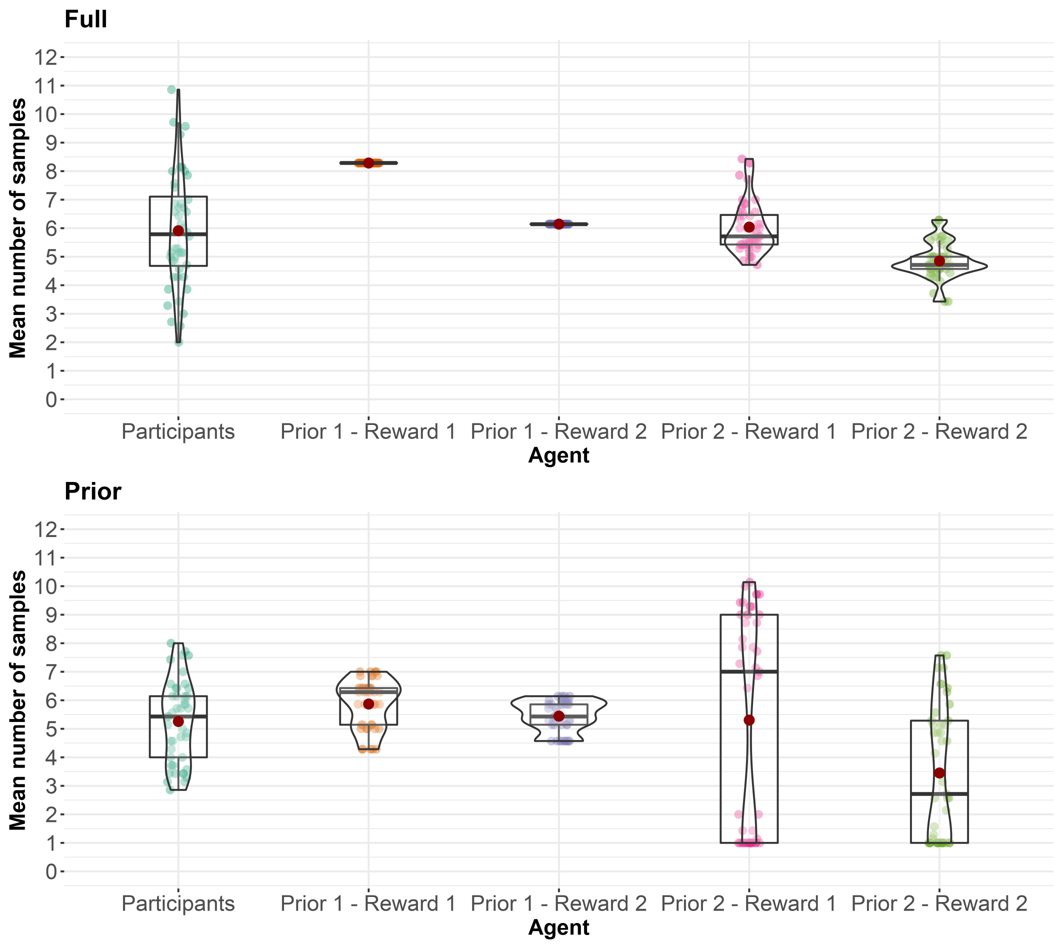

FIGURE S9: Model 1 (which uses a prior defined by objective values: Prior 1) assigned reward values only to the top three ranking options in the sequence (Reward 1), while Model 2 (which uses a prior defined by subjective values: Prior 2) assigned reward values commensurate with the subjective value of an option (Reward 2). To delineate whether the reward function of the optimal model can affect sampling biases, we have here plotted the distributions of the mean number of samples for participants versus Prior 1 with reward functions Reward 1 and Reward 2, and Prior 2 with reward functions Reward 1 and Reward 2, for the *full* and *prior* conditions. To reiterate, in the main text Prior 1 - Reward 1 is Model 1 for the *prior* condition, and Prior 2 - Reward 2 is Model 2 for the *full* condition. The red dots represent the mean, horizontal black lines represent the median, boxes show the 25% and 75% quantiles, and the whiskers represent the 95% confidence intervals.

Chapter 6

Neural computations of prospective social decisions

Abstract

An optimal stopping problem is a specific type of sequential decision-making scenario that requires people to make prospective decisions. Specifically, a decision maker must decide either to take an option now or to decline with the prospective expectation of a better future option, under the restriction that declined options cannot be returned to. A robust decision network has been found to underlie *take* versus *decline* decisions on (economic) sequential information sampling tasks where participants sample too few options compared to computational models of optimality. Here, we investigated whether a similar pattern of activation is associated with oversampling biases on a social optimal stopping task where participants attempted to maximise the facial attractiveness of a prospective date. Our results indicate that decisions to *take* versus *decline* an option engaged a network that closely overlaps the network identified in previous research, including the prefrontal cortex, parietal cortex, insula, and striatum. As such, our findings showed that activation in this decision network extends from previous economic optimal stopping tasks where people sample too few options, to ecologically-valid prospective social decisions where people sample too many options.

1 Introduction

In many real-world decision scenarios, a decision maker (DM) is required to make a trade-off: weighing a current reward against the prospective probability of a future reward. This is called prospective decision-making (Kolling et al., 2018). Optimal stopping problems are one example of a decision scenario that requires prospective decision-making (for historical reviews, see Ferguson, 1989; Freeman, 1983). Everyday optimal stopping problems that a DM may encounter are finding a parking space (Todd & Gigerenzer, 2012), renting an apartment (Zwick et al., 2003), buying a secondhand car (Costa & Averbeck, 2015), or choosing a date (Furl et al., 2019). What defines an optimal stopping problem is that options are sequentially presented in a fixed length sequence, and a DM must decide either to take an option now or to decline with the prospective expectation of a better future option. The difficult trade-off between taking and declining an option can be solved using a finite horizon Bayesian Markov decision process (MDP) (Costa & Averbeck, 2015).

Previous research on economic optimal stopping problems has revealed a decision-making network comprising the parietal and dorsolateral prefrontal cortices, ventral striatum, anterior insula, and anterior cingulate which could underlie the decision to stop sampling and choose an option (Costa & Averbeck, 2015). This finding was similar to earlier findings of Furl and Averbeck (2011) on the beads task (a related information sampling task, see Chapter 1, Sidebar 2). On the beads task, the decision to stop sampling was related to activation in the anterior insula, anterior cingulate, dorsal parietal cortex, and ventral striatum (Furl & Averbeck, 2011). Only activation in the dorsolateral prefrontal cortex was unique to Costa and Averbeck (2015). The findings of these two studies indicate that there is a consistent pattern of activation related to the belief that the expected reward for the current option exceeds the prospective reward of future options, motivating the decision to choose the current option. Aside from the functional magnetic resonance imaging (fMRI) results, Furl and Averbeck (2011) and Costa and Averbeck (2015) also found similar behavioural results. That is, compared to a Bayesian ideal observer model, participants sampled less evidence than optimal (*undersampling*). As such, the decision network identified by Furl and Averbeck (2011) and Costa and Averbeck (2015) could be specifically related to decisions to stop sampling too early on optimal stopping tasks.

Human undersampling biases are a common finding in the optimal stopping literature (e.g., Bearden et al., 2006; Cardinale et al., 2021; Guan & Lee, 2018; Hauser et al., 2018; Seale & Rapoport, 1997). However, there are also reports of participants sampling too much before taking an option (*oversampling*). Furl et al. (2019), for example, found evidence of oversampling on a social decision scenario where participants had to choose the most attractive face from a sequence of faces as their date. To the best of our knowledge, the neural network underlying oversampling biases on prospective social decisions is yet to be investigated. Therefore, the aim of the current study is to test whether the same decision network that contributed to the decision to take an option in Costa and Averbeck (2015) and Furl and Averbeck (2011) is involved in prospective social decisions on the facial attractiveness task (Furl et al., 2019). If we find activation consistent with previous findings, this evidence will bolster support for the general role of these regions in prospective decision-making, including difficult, ecologically-valid social decisions.

2 Materials and methods

2.1 Participants

Thirty neurotypical volunteers were included in our study (8 male and 22 female participants, M_{age} = 20.07 years, SD_{age} = 1.76 years). Participants were recruited through Royal Holloway, University of London's (RHUL) online Psychology Experiment Management System (paid pool) and Facebook groups intended for students of RHUL. All participants had normal or corrected-to-normal vision and identified as heterosexual. The study was approved by Royal Holloway, University of London's Ethics Board.

2.2 Experimental design

Our experimental design closely follows that of the facial attractiveness paradigm described by Furl et al. (2019). The experiment consisted of two parts: a rating task, which was carried out on a computer in the lab, and an optimal stopping task, which was carried out in the fMRI scanner. As our sample included only heterosexual participants, all male participants rated female faces, and all female participants rated male faces.

In the first part of the experiment, participants were instructed to rate 426 unique faces on their attractiveness, using a scale from 1 (very unattractive) to 9 (very attractive). Participants responded by keyboard press. The set of 426 faces used in our experiment was the same set used in Study 2 of Furl et al. (2019), and showed youthful individuals with happy expressions, roughly aged 18 to 30, ranging in viewpoint degree between frontal and three-quarter view, in colour with a circular grey mask. Each face was rated three times, allowing for a final attractiveness rating for each face to be computed from the mean of these ratings. Because of ongoing restrictions related to COVID-19, it was necessary to use MATLAB (MATLAB, 2015) for the delivery of the rating phase to the first 15 participants, and a Gorilla Experiment Builder (Anwyl-Irvine et al., 2020) version with only slightly different aesthetics for the remainder.

In the second part of the experiment, participants carried out an optimal stopping task in which they had to choose the most attractive face from sequences of faces as their date. Participants were told that the faces they would encounter in part two would be a random subset of the faces in part one, but were not provided with any further information regarding the attractiveness distribution that generated the sequences. The task was divided into four sessions, each of which contained eight sequences of eight faces, i.e., participants were shown 32 sequences in total. In line with the typical characteristics of optimal stopping problems, participants were unable to return to a previously rejected option, and they had no way of knowing for certain the value of the option(s) yet to come in any given sequence. The number of options remaining was shown at the top of the screen, and the rejected options were shown at the bottom of the screen. If participants did not choose any of the faces in the sequence, the last option automatically became their chosen face by default. Once participants had chosen an option, they were directed to a feedback screen which displayed their chosen face and the text 'Here is your new date!'. Participants provided responses throughout the second part of the experiment using a button box. Stimulus timings and inter-trial intervals are shown schematically in Figure 6.1.

Participants' sampling behaviour on the optimal stopping task was compared to a Bayesian ideal observer model (Costa & Averbeck, 2015) where performance is Bayesian optimal. Mathematically, the model is based on a discrete time, finite-horizon MDP with continuous states (for a full mathematical description we refer to Furl et al. (2019) and Costa and Averbeck (2015) [1]). Model values were fixed to the values used by Furl et al. (2019). For each option in the sequence, the ideal observer model computes the prospective reward values for taking and declining that option, and chooses the action with the highest value. The model received as input the same option values as the participants, in the order in which they were presented

[1] A mathematical description of the Bayesian ideal observer model is also provided in Chapters 3, 4 and 5 of this thesis.

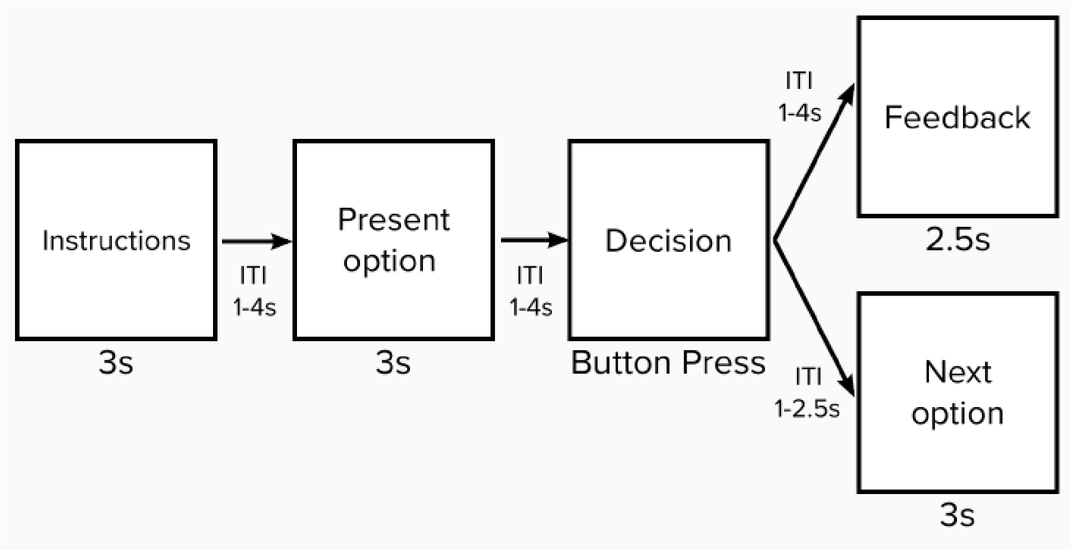

FIGURE 6.1: Schematic overview of the stimulus timings and inter-trial intervals of the optimal stopping task.

to the participants. Option values comprised the mean attractiveness rating for that particular option. Each participants' attractiveness ratings were log transformed before being put into the model to approximate normality, which helps satisfy the model's assumption that the generating distribution is normal. The cost-to-sample parameter was fixed to zero.

2.3 fMRI scanning

For our structural MRI and fMRI measurements, we used a Siemens Magnetom TrioTim syngo MR B17 3-T scanner (Siemens AG, Munich, Germany) at the CUBIC imaging centre at Royal Holloway, University of London, and an 8-channel head coil. We carried out a standard T1-weighted 1x1x1mm whole-brain structural scan (3D MPRAGE) with 1900ms time to repeat (TR), 3.03ms echo time (TE), flip angle 11°, matrix: 256 x 256mm. Echo-planar volumes were collected using a high-resolution 2x2x2mm multi-band sequence. Volumes were collected as 56 slices, 2.0mm thick, in-plane resolution 1200ms TR, 36.8ms TE, flip angle 30°. The first five 'dummy scans' of each session were discarded to allow for magnetization equilibration effects.

2.4 fMRI data preprocessing and statistical analysis

Data were preprocessed and analysed using MATLAB and SPM12 software (Wellcome Trust Centre for Neuroimaging, London; http://www.fil.ion.ucl.ac.uk/spm). All participants were retained after motion correction, as motion estimates (R

= [X Y Z pitch yaw roll]) were within reasonable bounds (maximum values of 8mm and 7 degrees). Motion parameters were put in as additional regressors in the first level analysis (explained in more detail below). Scans were realigned, spatially normalized to the Montreal Neurological Institute (MNI[2]) standardized space and smoothed to 5mm^3 full width half maximum using conventional procedures. Similar to Costa and Averbeck (2015), we employed a whole brain mass-univariate analysis for estimating the magnitude of the hemodynamic blood oxygenation level dependent (BOLD) response to each stimulus event (Friston et al., 1998).

At the individual-participant level, we computed 'first level' mass-univariate time series models for each participant. The aim of the first-level model is to predict the fMRI time series using regressors constructed from stick functions representing stimulus onset times, convolved with a canonical BOLD response function, which captures the temporal profile of the BOLD response to a behavioural event. Regressors included in the first level analysis were the motion parameters mentioned above, as well as the stick functions for the take events and decline events which were subtracted at the first level and those subtractions brought to the second level to be tested for significance from zero in a one sample t-test. For each regressor and its associated event type, a β-value was computed to estimate the magnitude of the BOLD response evoked by the events. β-values for different event regressors were then contrasted statistically to test whether the BOLD responses evoked by different events significantly differed in magnitude. Our first-level analyses also used an AR(1) autocorrelation model and a high-pass filter of 128 s. The term β-value or β-weight comes from the fact that it is a coefficient from a general linear model.

Our primary hypothesis concerned BOLD responses related to the decision to take the current option. Therefore, the participant response was modelled separately for choices to take the current option versus choices to decline the current option, so we could subsequently contrast these two decisions (*take* $>$ *decline* and *decline* $>$ *take*). Once β-values had been computed for both of our events of interest, as well as contrasts over these β-values (see Results), we then brought these β-values to a 'second level' analysis where they were entered into one-sample t-tests, treating participants as a random effect. This allowed us to test whether β-values were statistically consistent

[2]MNI is the standard template for fMRI research (Chau & McIntosh, 2005; Poldrack et al., 2011)

across participants. All results reported below were observed at a $p < .001$ uncorrected cluster detection threshold with a minimum of 100 voxels, after which they were tested for family-wise error (FWE) correction at the cluster level at $P(\text{FWE}) < .05$ using Gaussian random field theory. This is a conventional method (Penny et al., 2011) that uses the estimated smoothness of the data to correct for the massive number of multiple comparisons at all voxels in the whole brain.

3 Results

3.1 Behavioural results

After comparing participants' sampling behaviour to that of a Bayesian ideal observer model, we found that participants sampled more than the model ($Z = 0.501$, $p < .05$ [3]), and ended up with lower-ranked faces ($t(29) = -9.733$, $p < .001$) (Figure 6.2).

3.2 Functional imaging results

First, we examined the contrast between choices to take the current option versus choices to decline the current option. This contrast revealed activation in a number of areas (Table 6.1), including the dorsomedial prefrontal cortex (Figure 6.3A), extending to the anterior cingulate (Figure 6.3A), striatum (Figure 6.3A), insula (Figure 6.3B), and parietal cortex (Figure 6.3C). Like Furl et al. (2019), we found effects in the right supramarginal gyrus (an area involved in emotional processing; Silani et al., 2013). Additionally, we found effects in the occipital gyrus and the temporal gyrus (areas both related to the processing of faces; Jacques et al., 2019). The opposite contrast, declining versus taking the current option, revealed effects in areas related to visual processing (fusiform gyrus) and movement (supplementary motor area, postcentral gyrus, precentral gyrus). Furl and Averbeck (2011) also reported activity in the superior frontal gyrus, but we found activity in the middle/inferior frontal gyrus. Both contrasts were associated with activity in the cerebellum ansiform lobule, which is an area associated with cognitive and visuomotor functions (Sugihara, 2018).

[3] A Shapiro-Wilk test of normality indicated that the distribution of the mean number of samples for the optimal model differed significantly from the normal distribution ($W = 0.908$, $p = .013$), which is why a Wilcoxon signed rank test was used to calculate the difference between participants and the model for the mean number of samples

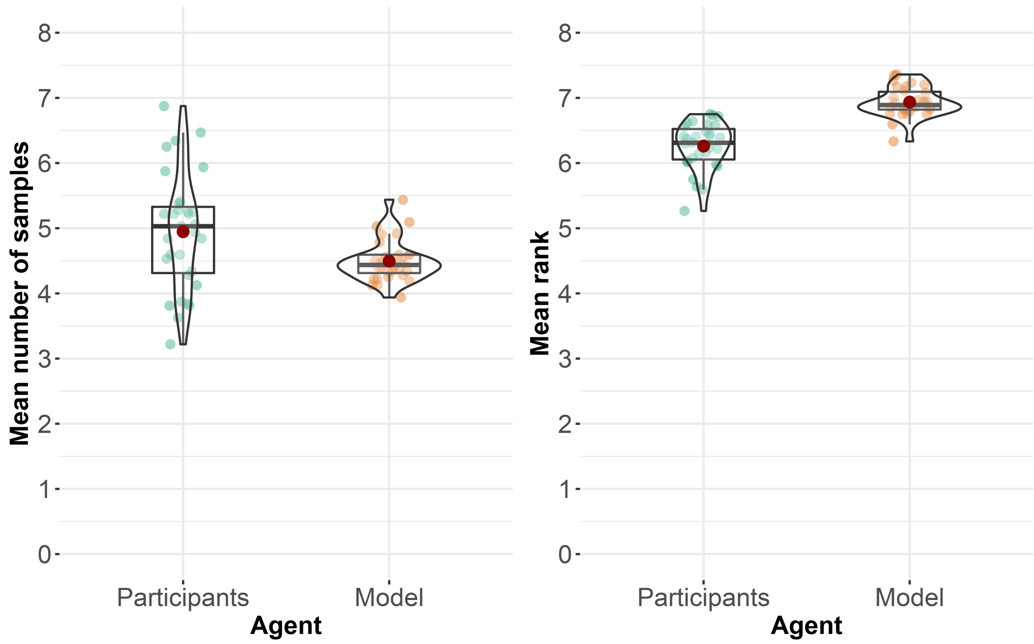

FIGURE 6.2: On the left: density plots of the mean number of samples, separately for each agent, overlaid with the individual data points and a box plot showing the mean and 95% confidence interval. On the right: density plots of the mean rank of the chosen option, separately for each agent, overlaid with the individual data points and a box plot showing the mean and 95% confidence interval.

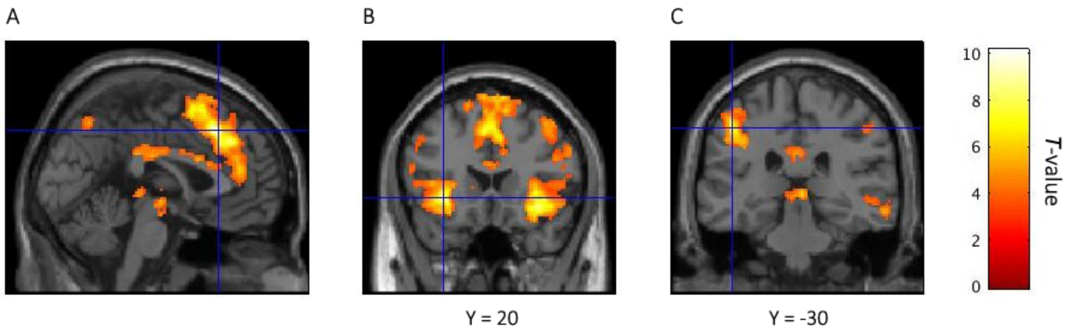

FIGURE 6.3: *Take* versus *decline* activation. A: prefrontal activation. B: insula activation. C: parietal activation.

4 Discussion

The fMRI study described in this paper investigated which neural networks underlie oversampling biases on a social optimal stopping task. In line with our expectations, participants sampled more options than a Bayesian ideal observer model. This finding is consistent with previous results on the so-called facial attractiveness task (Chapter 4; Furl et al., 2019). When comparing decisions to take an option with the decision to decline an option and

TABLE 6.1: Peak activation coordinates and statistics for areas differentiating take and decline choices.

Region	x^a	y	z	t^b	Size^c
Take > Decline					
Left frontal gyrus (medial/superior)	0	26	44	10.17	8026
Left insula	-32	20	-4	6.2	1655
Left postcentral gyrus	-46	-30	46	7.68	2612
Right supramarginal gyrus	46	-36	46	4.98	1347
Left inferior occipital gyrus	-42	-72	-10	4.92	683
Left cerebellum ansiform lobule	-8	-72	-26	4.57	210
Right middle temporal gyrus	62	-32	-12	4.56	161
Right caudate nucleus	14	8	14	4.51	108
Left middle frontal gyrus	-30	50	10	5.4	271
Decline > Take					
Right fusiform gyrus	32	-78	-14	18.39	8022
Left middle frontal gyrus	-32	20	6	7.89	398
Left supplementary motor area	2	12	50	7.89	1375
Left superior parietal lobule	-20	-62	52	7.19	426
Right inferior frontal gyrus	42	32	18	5.32	309
	38	6	30	5.34	176
Left cerebellum ansiform lobule	-6	-72	-32	5.24	108
Right insula	34	18	0	5.22	517
Left postcentral gyrus	-44	-30	46	5.05	984
Right precentral gyrus	32	-2	50	4.61	310

[a] Peak coordinates are reported in MNI space.
[b] Whole brain corrected at $P < .05$ FWE.
[c] Cluster volume estimated using 2 mm^3 voxels.

continue sampling, we found a robust network of activation that included a large cluster in the prefrontal cortex, striatum, insula, and parietal cortex. This pattern of activation was similar to the activation in Costa and Averbeck (2015) and Furl and Averbeck (2011), who both investigated the *take* versus *decline* contrast on an information sampling task.

Our results confirm the role of the frontal-parietal areas as well as cingulate, insula and striatal areas in prospective decision-making, supporting the hypothesis that these areas comprise the main decision network for decisions to commit to an option and stop sampling more information. Background information on the functions of these regions is provided in Chapter 1, Section 1.5. The current study builds upon previous literature by showing that activation in the decision network is independent of sampling biases (undersampling/oversampling) and decision-making domains (economic/social). While we have identified a decision network that seems to relate to prospective decision making on optimal stopping tasks, the functions of the individual regions comprising the decision network remain speculative. Based on the findings of previous research, we hypothesise that activation in the parietal cortex is associated with decision thresholds and sampling rate. The fact that we also found activity in the parietal cortex in the *decline* $>$ *take* contrast is unsurprising in light of previous research suggesting that the parietal cortex is involved in the integration of evidence as it is collected (Kiani & Shadlen, 2009). Furthermore, activation in the insula could be related to the value of the current option (Furl & Averbeck, 2011), while activation in the striatum might be associated with reward magnitude (Costa & Averbeck, 2015; Haber & Knutson, 2010). The role of the prefrontal cortex may be more general as it is an important contributor to cognitive control and decision-making (Domenech & Koechlin, 2015; Ridderinkhof et al., 2004), although some areas of the prefrontal cortex, such as the VMPFC and the ACC, have been highlighted in previous research to have functions that are key also to decision making on optimal stopping tasks. For example, the ACC was found to be important for evaluating potential future actions in foraging tasks that require exploratory behaviours (Kolling et al., 2016; Schuck et al., 2015; Wittmann et al., 2016) and activation in the VMPFC was associated with the subjective value of a stimulus (Brosch & Sander, 2013; Levy & Glimcher, 2012). We encourage future research to investigate our post-hoc hypotheses and study the differential contributions of the areas comprising the decision network.

In addition to studying the individual contributions of brain areas to decision making in general, future research may wish to further explore what cognitive processes are unique to taking an option vs declining an option specifically on optimal stopping tasks. Importantly, the main interest of the current study in the take vs decline contrast was simply the extent to which it would match previous studies (Costa & Averbeck, 2015; Furl & Averbeck, 2011) - a question our study has satisfied. As such, we did not have any a priori hypotheses regarding what different areas are doing when taking or declining an option and are unable to draw direct inferences about computations in these brain regions on the basis of the take versus decline contrast alone. One way to resolve this would be to fit our models to the behavioural data, after which the computational qualities, such as the value of taking an option vs the value of declining an option which vary trial by trial, can be extracted. These computational qualities could then be correlated with the fMRI data. This procedure is similar to the analyses described in Costa and Averbeck (2015) and Furl and Averbeck (2011). In Costa and Averbeck (2015), for example, the authors used the value estimates generated by their MDP model to produce a parametric modulator of the value of declining the current option minus the value of taking the current option. They found that this parametric modulator was significant amongst others in the VMPFC, meaning that activity in the VMPFC corresponded to the magnitude of the difference in value between take vs decline. Nevertheless, the primary goal of the current study was to show concordance with other studies, to which we have succeeded.

A limitation of the present study is that there also appeared to be some noise in the data, including peak activation coordinates not corresponding to brain areas in MNI space (not further reported nor included in Table 6.1). Although we have implemented standardised procedures for reducing noise in the data, differentiating signal from noise can be challenging (Liu, 2016). Also of note is that noise in the form of typical spatial patterns has been reported for studies using multi-band sequencing due to the simultaneous acquisition of multiple slices (Griffanti et al., 2017). Such artefacts were not observed in the current dataset, perhaps because they typically depend on the head movements of the participants (Griffanti et al., 2017), but we did not correct for multi-band scanning either. As such, we cannot exclude that noise due to the multi-band sequencing procedure was present in the data. Peak coordinates in regions of MNI space outside of the brain have been treated as noise, but we make no assumptions regarding our reported findings. However, activation in the frontal-parietal areas and insula in the *take* > *decline* contrast survived a more conservative correction method (whole-brain FWE at $p < .001$). This finding bolsters support for the role of these areas in prospective decisions.

Another limitation of the current paradigm is the fact that decline events were always followed by another option, whilst take events were always followed by the feedback screen that communicated the reward value of the chosen option. Because of this, some of the network activity observed could have been related to reward anticipation, rather than prospective decision-making. Regions typically involved in reward anticipation include the orbitofrontal cortex (Gorka et al., 2015; Kahnt et al., 2010), ventral striatum (Haber & Knutson, 2010; Jauhar et al., 2021; Knutson et al., 2001a; Knutson et al., 2001b), and anterior cingulate cortex (Gorka et al., 2015; Wilson et al., 2018). Although there is some activation in areas related to reward anticipation in the *take* $>$ *decline* contrast (e.g., right caudate nucleus), the evidence is not overwhelming. An explanation for this could be that the intrinsic reward as implemented in our paradigm was insufficient to elicit a strong response in regions associated primarily with monetary reward anticipation. However, Costa and Averbeck (2015) did offer their participants monetary rewards but nonetheless reported a similar pattern of activation. This suggests that activation in the *take* $>$ *decline* contrast is predominantly related to prospective decision-making. Future research may wish to further investigate whether some activation in the *take* $>$ *decline* contrast could be explained by reward anticipation by directly comparing an optimal stopping task that implements a feedback screen and rewards (intrinsic or monetary) to an optimal stopping task without a feedback screen and rewards.

In conclusion, participants showed an oversampling bias compared to a Bayesian ideal observer model on a social optimal stopping task. Prospective decision-making on this task was linked to activation in a decision network comprising the prefrontal cortex, insula, striatum, and the parietal cortex. Our findings indicate that activation in this decision network extends from economic optimal stopping tasks where people sample too few options, to ecologically-valid social decisions where people sample too many options.

References

Anwyl-Irvine, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2020). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods*, *52*(1), 388–407. https://doi.org/10.3758/s13428-019-01237-x

Bearden, J. N., Rapoport, A., & Murphy, R. (2006). Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Science*, *52*(9), 1437–1449. https://doi.org/10.1287/mnsc.1060.0535

Brosch, T., & Sander, D. (2013). Neurocognitive mechanisms underlying value-based decision-making: From core values to economic value. *Frontiers in Human Neuroscience*, *7*.

Cardinale, E. M., Pagliaccio, D., Swetlitz, C., Grassie, H., Abend, R., Costa, V. D., Averbeck, B. B., Brotman, M. A., Pine, D. S., Leibenluft, E., & Kircanski, K. (2021). Deliberative Choice Strategies in Youths: Relevance to Transdiagnostic Anxiety Symptoms. *Clinical Psychological Science*, 1–11. https://doi.org/10.1177/2167702621991805

Chau, W., & McIntosh, A. R. (2005). The Talairach coordinate of a point in the MNI space: How to interpret it. *Neuroimage*, *25*(2), 408–416. https://doi.org/10.1016/j.neuroimage.2004.12.007

Costa, V. D, & Averbeck, B. B. (2015). Frontal-Parietal and Limbic-Striatal Activity Underlies Information Sampling in the Best Choice Problem. *Cerebral Cortex*, *25*(4), 972–982. https://doi.org/10.1093/cercor/bht286

Domenech, P., & Koechlin, E. (2015). Executive control and decision-making in the prefrontal cortex. *Current Opinion in Behavioral Sciences*, *1*, 101–106. https://doi.org/10.1016/j.cobeha.2014.10.007

Ferguson, T. S. (1989). Who Solved the Secretary Problem? *Statistical Science*, *4*(3), 282–289. https://doi.org/10.1214/ss/1177012493

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

Friston, K., Fletcher, P., Josephs, O., Holmes, A., Rugg, M., & Turner, R. (1998). Event-Related fMRI: Characterizing Differential Responses. *Neuroimage*, *7*(1), 30–40. https://doi.org/10.1006/nimg.1997.0306

Furl, N., & Averbeck, B. B. (2011). Parietal Cortex and Insula Relate to Evidence Seeking Relevant to Reward-Related Decisions. *Journal of Neuroscience*, *31*(48), 17572–17582. https://doi.org/10.1523/JNEUROSCI.4236-11.2011

Furl, N., Averbeck, B. B., & McKay, R. T. (2019). Looking for Mr(s) Right: Decision bias can prevent us from finding the most attractive face. *Cognitive psychology*, *111*, 1–14. https://doi.org/10.1016/j.cogpsych.2019.02.002

Gorka, S. M., Phan, K. L., & Shankman, S. A. (2015). Convergence of EEG and fMRI measures of reward anticipation. *Biological Psychology*, *112*, 12–19. https://doi.org/10.1016/j.biopsycho.2015.09.007

Griffanti, L., Douaud, G., Bijsterbosch, J., Evangelisti, S., Alfaro-Almagro, F., Glasser, M. F., Duff, E. P., Fitzgibbon, S., Westphal, R., Carone, D., Beckmann, C. F., & Smith, S. M. (2017). Hand classification of fMRI ICA noise components. *NeuroImage*, *154*, 188–205. https://doi.org/10.1016/j.neuroimage.2016.12.036

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339–361. https://doi.org/10.1037/dec0000081

Haber, S. N., & Knutson, B. (2010). The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology*, *35*(1), 4–26. https://doi.org/10.1038/npp.2009.129

Hauser, T. U., Moutoussis, M., Purg, N., Dayan, P., & Dolan, R. J. (2018). Beta-Blocker Propranolol Modulates Decision Urgency During Sequential Information Gathering. *The Journal of Neuroscience*, *38*(32), 7170–7178. https://doi.org/10.1523/JNEUROSCI.0192-18.2018

Jacques, C., Jonas, J., Maillard, L., Colnat-Coulbois, S., Koessler, L., & Rossion, B. (2019). The inferior occipital gyrus is a major cortical source of the face-evoked N170: Evidence from simultaneous scalp and intracerebral human recordings. *Human Brain Mapping*, *40*(5), 1403–1418. https://doi.org/10.1002/hbm.24455

Jauhar, S., Fortea, L., Solanes, A., Albajes-Eizagirre, A., McKenna, P. J., & Radua, J. (2021). Brain activations associated with anticipation and delivery of monetary reward: A systematic review and meta-analysis of fMRI studies. *PloS one*, *16*(8). https://doi.org/10.1371/journal.pone.0255292

Kahnt, T., Heinzle, J., Park, S. Q., & Haynes, J.-D. (2010). The neural code of reward anticipation in human orbitofrontal cortex. *Proceedings of the*

National Academy of Sciences, *107*(13), 6010–6015. https://doi.org/10.1073/pnas.0912838107

Kiani, R., & Shadlen, M. N. (2009). Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science*, *324*(5928), 759–764. https://doi.org/10.1126/science.1169405

Knutson, B., Adams, C. M., Fong, G. W., & Hommer, D. (2001a). Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens. *Journal of Neuroscience*, *21*(16). https://doi.org/10.1523/JNEUROSCI.21-16-j0002.2001

Knutson, B., Fong, G. W., Adams, C. M., Varner, J. L., & Hommer, D. (2001b). Dissociation of reward anticipation and outcome with event-related fMRI. *Neuroreport*, *12*(17), 3683–3687.

Kolling, N., Scholl, J., Chekroud, A., Trier, H. A., & Rushworth, M. F. S. (2018). Prospection, Perseverance, and Insight in Sequential Behavior. *Neuron*, *99*(5), 1069–1082. https://doi.org/10.1016/j.neuron.2018.08.018

Kolling, N., Wittmann, M. K., Behrens, T. E. J., Boorman, E. D., Mars, R. B., & Rushworth, M. F. S. (2016). Value, search, persistence and model updating in anterior cingulate cortex. *Nature Neuroscience*, *19*(10), 1280–1285. https://doi.org/10.1038/nn.4382

Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, *22*(6), 1027–1038. https://doi.org/10.1016/j.conb.2012.06.001

Liu, T. T. (2016). Noise contributions to the fMRI signal: An overview. *Neuroimage*, *143*, 141–151. https://doi.org/10.1016/j.neuroimage.2016.09.008

MATLAB. (2015). Version 2015b. The MathWorks Inc. *Natick, Massachusetts*.

Penny, W. D., Friston, K. J., Ashburner, J. T., Kiebel, S. J., & Nichols, T. E. (2011). *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Elsevier.

Poldrack, R. A., Mumford, J. A., & Nichols, T. E. (2011). *Handbook of Functional MRI Data Analysis*. Cambridge University Press.

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The Role of the Medial Frontal Cortex in Cognitive Control. *Science*, *306*(5695), 443–447. https://doi.org/10.1126/science.1100301

Schuck, N. W., Gaschler, R., Wenke, D., Heinzle, J., Frensch, P. A., Haynes, J.-D., & Reverberi, C. (2015). Medial Prefrontal Cortex Predicts Internally Driven Strategy Shifts. *Neuron*, *86*(1), 331–340. https://doi.org/10.1016/j.neuron.2015.03.015

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

Silani, G., Lamm, C., Ruff, C. C., & Singer, T. (2013). Right Supramarginal Gyrus Is Crucial to Overcome Emotional Egocentricity Bias in Social Judgments. *Journal of Neuroscience*, *33*(39), 15466–15476.

Sugihara, I. (2018). Crus I in the Rodent Cerebellum: Its Homology to Crus I and II in the Primate Cerebellum and Its Anatomical Uniqueness Among Neighboring Lobules. *The Cerebellum*, *17*(1), 49–55. https://doi.org/10.1007/s12311-017-0911-4

Todd, P. M., & Gigerenzer, G. (2012). *Ecological Rationality Intelligence in the World*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195315448.001.0001

Wilson, R. P., Colizzi, M., Bossong, M. G., Allen, P., Kempton, M., Abe, N., Barros-Loscertales, A. R., Bayer, J., Beck, A., Bjork, J., Boecker, R., Bustamante, J. C., Choi, J. S., Delmonte, S., Dillon, D., Figee, M., Garavan, H., Hagele, C., Hermans, E. J., … MTAC. (2018). The Neural Substrate of Reward Anticipation in Health: A Meta-Analysis of fMRI Findings in the Monetary Incentive Delay Task. *Neuropsychology Review*, *28*(4), 496–506. https://doi.org/10.1007/s11065-018-9385-5

Wittmann, M. K., Kolling, N., Akaishi, R., Chau, B. K. H., Brown, J. W., Nelissen, N., & Rushworth, M. F. S. (2016). Predictive decision making driven by multiple time-linked reward representations in the anterior cingulate cortex. *Nature Communications*, *7*(1). https://doi.org/10.1038/ncomms12327

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer Sequential Search: Not Enough or Too Much? *Marketing Science*, *22*(4), 503–519. https://doi.org/10.1287/mksc.22.4.503.24909

Chapter 7

Discussion

In my introduction, I specified three sub-questions, namely: are sampling biases dependent on the decision-making domain? Can certain task features explain sampling biases? And, which brain areas correlate with prospective decision-making on optimal stopping tasks? In the following sections, I outline how the research included in this thesis has addressed these research questions, I discuss the theoretical contributions of this thesis as well as the implications for the field, and I suggest directions for future research.

1 Sub-question 1: sampling biases are not dependent on the decision-making domain

My first empirical chapter (Chapter 3) investigated whether oversampling biases on full information optimal stopping problems are an intrinsic feature of personal mate choice decisions, as suggested by Furl et al. (2019). In the first study, I showed that oversampling biases persisted when participants made decisions about the attractiveness of potential dates on an imaginary client's behalf (*matchmaker*). In the second study, I showed that oversampling biases persisted when participants made decisions about the trustworthiness of faces, rather than attractiveness (*trustworthiness*). As such, the studies described in Chapter 3 provide evidence that oversampling biases are not limited to the personal mate choice domain, and may therefore not be domain specific. However, this conclusion remained limited to decision-making domains that used images of faces as stimuli. Up to this point, all studies that found oversampling biases in participants (Furl et al., 2019, matchmaker, trustworthiness) had used images of faces. Although I show in Chapter 3 that the different paradigms lead to distinctly different judgements (Figure

3.1), it remains unclear whether oversampling biases extend to other image-based domains beside faces. This was one of the main objectives of my second empirical chapter (Chapter 4).

Chapter 4 builds upon the results of Chapter 3 by investigating whether sampling rates across different image-based decision-making domains depend on different sampling biases. The results of two studies provided convergent evidence that participants oversample in three different image-based decision-making domains: faces, food and holiday destinations, thus providing further support for the theory that sampling biases are not domain specific. Yet a limitation remained: to this point I had only shown oversampling biases on image-based decision-making domains. To be able to conclude that sampling biases are not dependent on the decision-making domain, I would need to show that participants can oversample on number-based domains. To this point, research on number-based full information problems had mainly reported that participants undersample on these kinds of tasks (Baumann et al., 2020; Cardinale et al., 2021; Costa & Averbeck, 2015). Hence, one of the objectives of Chapter 5 was to determine whether participants would continue to oversample on our full information task when numeric stimuli were used instead, but all other task features as implemented in Chapters 3 and 4 were maintained.

The results of the first two studies described in Chapter 5 showed that oversampling biases persisted when smartphone prices (i.e., numbers) were used as stimuli in the full information optimal stopping task. Why I observed oversampling biases on a number-based task while previous research reported undersampling is further discussed in Section 2. Considering the evidence presented in my first three experimental chapters, I am now able to answer my first research question and conclude that sampling biases are not dependent on the decision-making domain, whether image-based or number-based.

2 Sub-question 2: certain task features can explain sampling rates, but not sampling biases

Having ascertained that sampling biases (i.e., over- or undersampling compared to optimality) are not dependent on the decision-making domain, I now consider my second research question, whether other task features explain sampling biases. In Chapter 4, I investigated whether the moments of the generating distribution might explain the participants' and model's variations in sampling rate. Previous research had shown that on number-based full information problems, a more positively skewed generating distribution can increase participants' sampling rate (Baumann et al., 2020). Chapter 4 builds on the findings of Baumann et al. (2020) by investigating whether the moments of the generating distribution also affect participants' sampling rate on image-based full information problems, and whether the optimal model is affected in a similar way by the moments of the generating distribution.

At first glance, this hypothesis may not seem related to task features. However, from the results of Chapter 3 we know that different decision domains that use the same stimuli can lead to distinctly different generating distributions (Figure 3.1). Therefore, I hypothesised that using different images altogether like in Chapter 4 (faces, food and holiday destinations) could in a similar way also lead to differences in the moments of the generating distribution between conditions. As shown in Figures 4.3 and 4.4, this was indeed the case. Significant differences were found between conditions in the mean, variance, skewness and kurtosis of the generating distributions. Additionally, the results showed that the mean and skewness of the generating distribution could statistically explain sampling behaviour, as these moments correlated in a similar way with both participants' sampling rate and the optimal model's sampling rate. Moreover, the variance of the generating distribution showed a negative correlation with the optimal model's sampling rate, and the kurtosis showed a positive correlation with the model's sampling rate. As such, I demonstrate that using different images, or even just different judgements of the same image (i.e., attractiveness versus trustworthiness), can lead to variations in the shape of the generating distribution, which consequently affects sampling rate. However, the moments of the generating distribution cannot explain sampling biases because the optimal model's sampling rate is affected in a similar way to participants by variations in the shape of the generating distribution.

Nevertheless, there are additional task features that can be considered candidates for explaining sampling biases and that could therefore answer my second research question. In Chapter 5, after having found that the decision-making domain does not affect sampling biases, I deduced that some task feature(s) inherent to my full information task, other than the use of images, can lead to oversampling. To investigate this, I conducted a third study in Chapter 5 where I explored the effects of four different task features on participants' sampling biases. The task features in question were the so-called grey squares, the bonus payments, the fixed or self-paced timing of the task, and the rating phase. Yet, none of the task features was sufficient to instigate an oversampling bias, and no significant differences in participants' sampling rate were found across conditions. Instead, what caused sampling biases to differ was the behaviour of the optimal model. The model sampled less in the condition that used a slightly different specification of the generating distribution and reward function compared to the other five conditions, causing an oversampling bias to be observed in participants. The different model specifications and their implications are further discussed below in Section 4. For now, what is important is that I have found evidence that the optimal model, rather than certain task features, can explain sampling biases.

3 Sub-question 3: a decision network correlates with prospective decision-making on optimal stopping tasks

In my final empirical chapter, Chapter 6, I investigated which brain areas correlate with prospective decision-making on a full information optimal stopping task on which people show an oversampling bias. When comparing decisions to take an option versus decisions to decline an option (*take > decline*), a robust pattern of activation emerged which included the prefrontal cortex, striatum, insula, and parietal cortex. This pattern of activation was similar to the activation found in previous research (Costa & Averbeck, 2015; Furl & Averbeck, 2011), and confirms that activation in this decision network extends from a number-based full information task (where people undersample) to an image-based full information task (where people oversample). The functions of the individual regions comprising the decision network, however, remain speculative.

The activation in the prefrontal cortex was expected due to its role in cognitive control and decision-making (Domenech & Koechlin, 2015; Ridderinkhof et al., 2004). A hypothesis for future studies on the basis of the current findings and the literature is that specifically the DLPFC and VMPFC are important for decision-making on optimal stopping tasks due to their role in value computation (Costa & Averbeck, 2015; Lin et al., 2020) and evidence accumulation (Furl & Averbeck, 2011; Gluth et al., 2012). The striatum is know for its specific function in reward anticipation (Haber & Knutson, 2010; Knutson et al., 2001a; Knutson et al., 2001b). Costa and Averbeck (2015), for example, found that activation in the ventral striatum related to the size of the reward outcome. In my fMRI study I found striatal activation on a task that had no monetary rewards. Instead, I assumed that participants would feel intrinsically rewarded by the perceived quality of their choice. My findings suggest that this intrinsic reward was sufficient to instigate activation in the striatum. Alternatively, the activity I observed in the ventral striatum in the *take* versus *decline* contrast could have been related to the anticipation of the reward feedback screen which appeared after an option had been chosen. The activation I found in the insula bolsters the hypothesis proposed by Furl and Averbeck (2011) that on optimal stopping tasks, activation in the insula relates to the value of the current option. Finally, previous research has found that activation in the parietal cortex was associated with decision thresholds and sampling rate on optimal stopping tasks (Costa & Averbeck, 2015; Furl & Averbeck, 2011). I observed activation in the parietal cortex in both the *take* > *decline* and *decline* > *take* contrasts which suggests that on my full information problem as well, the parietal cortex is involved in the integration of evidence as it is collected (Kiani & Shadlen, 2009), thus supporting the findings of Costa and Averbeck (2015) and Furl and Averbeck (2011).

At this point, one might speculate why I observed a similar pattern of activation in a task where participants oversampled, compared to a task where participants undersampled (Costa & Averbeck, 2015). The answer is that sampling biases are relative to the sampling rate of the optimal model. From Figure 2E of Costa and Averbeck (2015), I can discern that participants on average sampled around 4.9 options on the task with sequence length eight, whilst the optimal model sampled around 5.1 options. On my optimal stopping task, which also had a sequence length of eight, participants on average sampled 4.9 options (SD = 0.87), whilst the optimal model on average sampled 4.49 option (SD = 0.32). In other words, the actual sampling rate of participants across the two studies was similar, explaining the strongly overlapping decision network, but the sampling rate of the model was different, leading to different results in sampling bias.

4 Theoretical contributions and implications for the field

This thesis focuses on a specific and simple version of a full information optimal stopping problem where the actual values of the options are presented, the distributions that generate the option values are known, there is no extrinsic (e.g., monetary) cost-to-sample, there is no recall of rejected options, and in most cases decision outcomes provide a reward equal to their value (Abdelaziz & Krichen, 2006; Cardinale et al., 2021; Costa & Averbeck, 2015; Furl et al., 2019; Gilbert & Mosteller, 1966; Guan et al., 2014; Hill, 2009; Lee, 2006; Shu, 2008). This is important to specify as many other versions of optimal stopping problems exist, each with their own optimal solution, and each (potentially) affecting human sampling behaviour in different ways (see Chapter 1 for a review). The full information problem, however, is a more novel area of investigation and, importantly, better resembles real-life optimal stopping scenarios than versions which operate under many strong theoretical assumptions. For example, earlier research has long focused on the rank-based version (i.e., the classical secretary problem) which presents participants with the relative rank of an option rather than the actual value of the option (Ferguson, 1989). However, in a real-life optimal stopping scenario such as accepting a job offer, or finding a parking spot, a decision maker is unlikely to encounter the relative rank of an option. Furthermore, the full information problem is a far more difficult problem to solve computationally compared to the relatively simple optimal solutions that exist for secretary problems, and so how participants solve such problems is especially interesting and important to understand. Therefore, the theoretical contributions of this thesis apply to the specific full information problem detailed above, and should not simply be generalised to other versions of optimal stopping problems.

Together, my results indicate that oversampling biases might be more common than previously thought. To illustrate this, I conducted a meta-analysis of the studies included in this thesis, treating conditions as separate studies (Figure 7.1). The location of the study (online or in the lab) and the specification of the optimal model (Model 1 or Model 2[1], as described in

[1]Recall that for Model 1, the mean and variance of the prior of the generating distribution were set to those of the log transformed distribution of raw smartphone prices. Reward values were assigned proportionally to the top three ranking options. For Model 2, the mean and variance of the prior of the generating distribution were set to those of the log

Chapter 5) were included as moderators. The results showed that the specification of the optimal model was a significant moderator (β = 1.722, z = 4.723, p < .001), but the location was not (β = -0.682, z = -1.790, p = .073). The meta-analysis reveals that oversampling biases seem associated primarily with the Model 2 specification of the optimal model. Recall that Model 2 resembled the optimal model used by Furl et al. (2019), who also reported oversampling. Could sampling biases be explained by the optimal model, rather than the decision-making domain or certain task features as initially hypothesised? My results suggest this is the case: in Chapter 5 I demonstrate that sampling biases can flip from oversampling to undersampling, and from undersampling to no sampling bias, when the alternative model was used instead (i.e., Model 1 instead of Model 2 or Model 2 instead of Model 1; Figure S6). A similar observation can be made when comparing my fMRI study (Chapter 6) to Costa and Averbeck (2015), as I have done in Section 3. Participants' sampling rate was almost identical in both studies. However, Costa and Averbeck (2015) used an optimal model like Model 1, which sampled more options than participants, resulting in an undersampling bias, while I used an optimal model like Model 2, which sampled fewer options than participants resulting in an oversampling bias.

These findings highlight the importance of researchers clearly and correctly specifying the parameters of their optimal model. As Griffiths et al. (2012) state, the goal of Bayesian models is not to show that people are optimal, the goal is to characterise the problem people are solving and its ideal solution. This means that the Bayesian model should provide the optimal solution to the problem, with the problem being specified from the participants' perspective. In other words, both participants and the optimal model should be trying to solve the same problem (e.g., they are both trying to maximise the objective/subjective value of the options), even if the optimal model and participants use different computations to solve this common problem.

In the case of optimal stopping problems where the generating distribution is one of the model specifications, it is critical that this generating distribution reflects the distribution that participants operate on when making decisions on the task. Therefore I suggest that using subjective option values rather than objective option values might be a more accurate representation of participants' generating distribution. As I have shown in Chapter 5,

transformed distribution of subjective (rating) values. Reward values were commensurate with the subjective value of the chosen option.

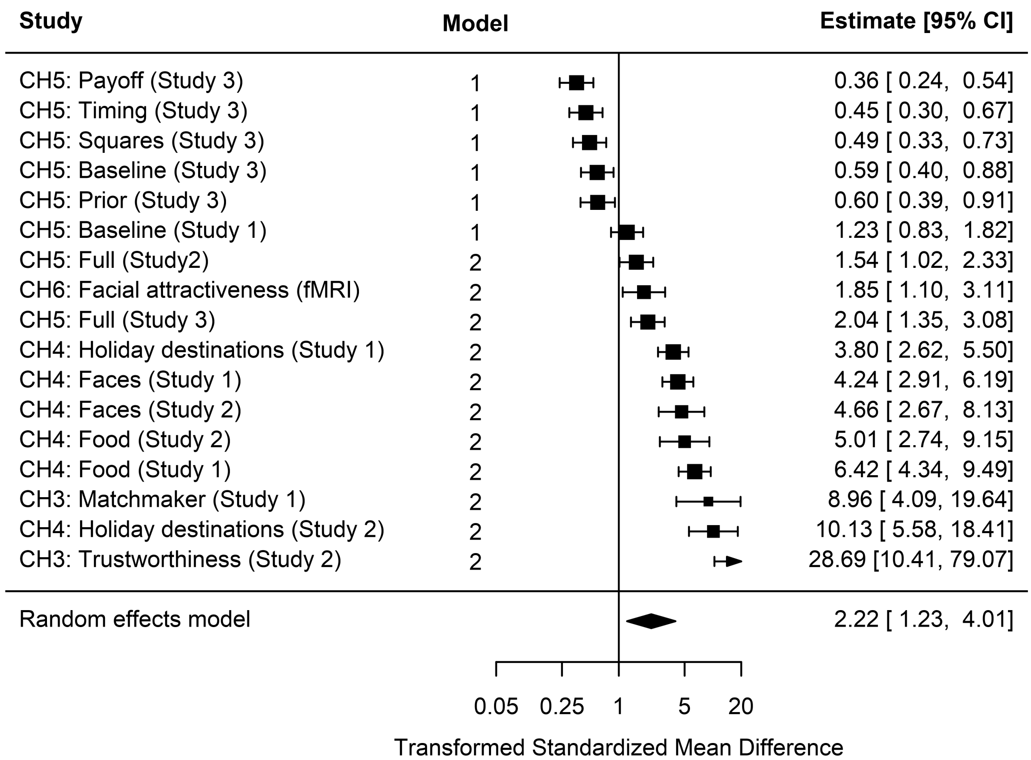

FIGURE 7.1: Forest plot showing log transformed standardized mean differences between participants and the optimal model for each of the conditions/studies included in this thesis. The analysis performed was a meta-analysis mixed effect model. The vertical line indicates the point on the x-axis equal to no effect. Squares represent the effect size of the study, with whiskers representing the 95 % confidence interval (CI). Studies with a larger weight have a larger square. The random effects model represents the average effect, with the length of the diamond shape symbolising the confidence interval of the pooled result.

options' relative ranks were preserved for the most part when using subjective values, but the shape of the distribution was different between subjective and objective values. And as I have illustrated in Chapter 4, the shape of the generating distribution can affect the sampling rate of both participants and the optimal model. As such I propose that using subjective values might be more accurate for both image-based and number-based tasks. One of the main benefits of using each participant's unique generating distribution of subjective values is that it ensures that the optimal model reflects the optimal solution for that particular participant.

The implication of the debate between using objective or subjective values for the generating distribution is that the model specification can differ from task to task, depending on the instructions and information provided to participants. This can be observed when comparing Costa and Averbeck (2015), Cardinale et al. (2021) and Furl et al. (2019), for example. Costa and Averbeck (2015) and Cardinale et al. (2021) used a similar Bayesian optimal model (like Model 1), and the findings of the two studies may therefore be compared (both reported undersampling). Furl et al. (2019), however, used a different specification of the generating distribution for their Bayesian optimal model (like Model 2), and found contrasting results (oversampling). Furl et al. (2019) then suggested that an explanation for the opposing findings could be the difference in decision-making domains. The theoretical contribution of this thesis is that research like Furl et al. (2019) cannot directly be compared to research like Costa and Averbeck (2015) and Cardinale et al. (2021), because the specification of the generating distribution for the optimal model is dissimilar. This highlights the importance of establishing standardised procedures for modelling optimality in full information problems to aid the comparison of optimal stopping research.

5 Limitations and future directions

It is important to recognise that individual differences are ever present in the study of human behaviour (Ozer, 1990), and the field of optimal stopping problems is no exception (Guan et al., 2014; Guan et al., 2015; Lee et al., 2005). An example of a study that reported strong individual differences is Sonnemans (1998). Thirty-six participants received 'bids' from the computer on an article they were supposed to sell for the highest possible price. The underlying distribution was known to participants and uniform, hence the problem could be classified as a full information problem. However, compared to the studies described in this thesis, there was a fixed cost-to-sample (two cents) and recall of previously rejected options was allowed. After the experiment, participants filled out a questionnaire about the strategies they used during the optimal stopping task. The results showed clear individual differences in recall (some participants never recalled, some recalled most of the time), and decision-making strategies. Next, Sonnemans (1998) grouped participants according to their decision-making strategy. Approximately two-thirds of participants could be classified as 'maximizers', while one-third of participants could be classified as 'satisficers'. Maximizers employed a strategy close to the optimal strategy, whereas satisficers focused mainly on earnings.

In general, participants sampled close to optimality, but with a tendency to stop sampling too early and to recall unnecessarily.

Supporting the idea of Sonnemans (1998) to cluster participants based on their decision-making strategy are the findings of Lee (2006) on a generalised secretary problem. In Lee (2006), participants were presented with 40 sequences of five options. Options were drawn from a uniform distribution of values between 0.00 and 100.00, and participants had to choose the highest value in a sequence (choosing any other value was wrong). Lee (2006) computed individual decision-making thresholds for 50 participants, superimposed these on the thresholds for the optimal decision rule (which decrease over sequence position), and revealed a clear variation across participants, although further clustering of participants was not attempted. A visual comparison, however, showed that some participants used thresholds close to the optimal thresholds, whilst others used either a single fixed threshold or varying thresholds that were non-optimal.

It is very likely that, like Sonnemans (1998) and Lee (2006), participants in my studies used various decision-making strategies. In Chapter 3, the results of the exploratory analysis indicated that three different models of human behaviour provided a good fit to the participants' sampling data: the sample reward model, the biased values model and the attractive prior model. One explanation for this is that human sampling behaviour can best be explained by different models in different participants. As such, my thesis, and especially the model comparison included in Chapter 3, provides some early hints at interesting individual differences which could motivate new lines of research.

To give an example, further research into whether certain individual traits can affect participants' sampling rate in a systematic way could be beneficial in the quest for factors that influence sampling biases on optimal stopping problems. Analytic cognitive style, for instance, has been found to predict data gathering on the beads task (a related information sampling task; Ross et al., 2016). Analytic cognitive style finds it origin in the dual-process theory of reasoning, and is defined as "the willingness or disposition to critically evaluate outputs from Type 1 processing and engage in effortful Type 2 processing" (p. 301, Ross et al., 2016), with Type 1 processing being defined as intuitive and automatic processing, and Type 2 processing being defined as slow and deliberative processing (Kahneman, 2011). As discussed in Chapter 6, the neural network underlying the decision to stop sampling was found to be almost identical for the beads task and an optimal stopping task (Costa & Averbeck, 2015; Furl & Averbeck, 2011). Therefore, it can be hypothesised that the findings regarding analytic cognitive style might also extrapolate to optimal stopping problems, where analytic cognitive style could predict individuals' sampling rate.

Additionally, as I have highlighted in Section 4, I believe it is important for future research to establish standardised procedures for (full information) optimal stopping tasks, as well as for modelling optimality on these tasks. For example, a set of guidelines could be established detailing how, in cases where subjective option values are not used, participants should be taught the objective generating distribution of option values. This topic has been addressed previously in the discussions of Chapter 4 and Chapter 5. Previous research has used many different approaches to teach participants the generating distribution of objective option values, including visual presentations of probability distributions (Baumann et al., 2020), statistical terminology (Guan et al., 2014), enriched feedback and/or financial rewards (Campbell & Lee, 2006), and repeated play (Goldstein et al., 2017, 2020). Others simply assumed that participants would be familiar with the real-world distribution of values that options were sampled from (Cardinale et al., 2021; Costa & Averbeck, 2015). And of course, the so-called rating phase as implemented in Study 3 of Chapter 5 (*prior* condition) could also be used for this purpose. However, it remains unclear which (if any) of these approaches best approximates the generating distribution participants are using to solve the optimal stopping problem. Based on the results of this thesis, I suggest that researchers may wish to use subjective option values instead as they might be more accurate at representing people's beliefs and they can account for individual differences.

6 Conclusion

The aim of this thesis was to explain human oversampling biases on full information optimal stopping problems using a variety of techniques including behavioural studies, computational analyses and neuroimaging methodology. First, I showed that sampling biases are not dependent on the decision-making domain. In a novel contribution to the literature, I demonstrate that oversampling biases extend from the mate choice domain to other decision-making domains such as trustworthiness, food and holiday destinations, as well as number-based domains such as smartphone prices. Second, I investigated an array of task features to determine whether they could explain sampling biases, but while I found that the moments of the generating distribution affected participants' sampling rate, none of the investigated task features were found to influence sampling biases. Finally, I have presented neuroimaging evidence indicating that similar areas in the so-called decision network are activated when a decision maker samples too few or too many options on a full information problem.

An interesting finding that emerged from my research was that the optimal model, and particularly the specification of the generating distribution, might be able to explain oversampling biases on full information problems. This was not one of my a priori hypotheses, but after conducting a meta-analysis of the studies included in this thesis, strong evidence for this hypothesis emerged. I investigate two different specifications of the prior distribution in this thesis, and show that the first specification (so-called Model 1) mostly led to undersampling, while the second specification (so-called Model 2) mostly led to oversampling. This highlights the importance of researchers clearly and correctly specifying the generating distribution. Particularly in the case of Bayesian optimal models, the idea is that the model characterises the problem people are solving, meaning that the generating distribution of the optimal model should correspond to the generating distribution that participants are using to solve the optimal stopping problem. Therefore, a key future direction for research is to determine how to best approximate participants' generating distribution. Based on my findings, I suggest that using subjective option values corresponding to each participant's own belief might be more accurate than using the objective option values. When using objective values, they must be learned by the participant in an appropriate way or they might remain an inaccurate representation of the participant's prior beliefs.

If there is one conclusion that could be drawn from my thesis and the optimal stopping literature it is that sampling rates, and indeed sampling biases, are heavily task-dependent. For example, in this thesis I have shown that the decision-making domain can affect participants' generating distribution of subjective values, which consequently leads to variations in sampling rates. I therefore encourage future research to work on establishing standardised procedures for (full information) optimal stopping tasks, as well for modelling optimality (see previous point). Not only would this benefit the extrapolation of scientific findings to real-world optimal stopping scenarios (e.g., accepting a job offer, finding a parking spot, swiping on Tinder), it would also aid the comparison of optimal stopping research.

References

Abdelaziz, F. B., & Krichen, S. (2006). Optimal stopping problems by two or more decision makers: A survey. *Computational Management Science*, *4*(2), 89–111. https://doi.org/10.1007/s10287-006-0029-5

Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23), 12750–12755. https://doi.org/10.1073/pnas.2002312117

Campbell, J., & Lee, M. D. (2006). The Effect of Feedback and Financial Reward on Human Performance Solving 'Secretary' Problems. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *28*(28).

Cardinale, E. M., Pagliaccio, D., Swetlitz, C., Grassie, H., Abend, R., Costa, V. D., Averbeck, B. B., Brotman, M. A., Pine, D. S., Leibenluft, E., & Kircanski, K. (2021). Deliberative Choice Strategies in Youths: Relevance to Transdiagnostic Anxiety Symptoms. *Clinical Psychological Science*, 1–11. https://doi.org/10.1177/2167702621991805

Costa, V. D, & Averbeck, B. B. (2015). Frontal-Parietal and Limbic-Striatal Activity Underlies Information Sampling in the Best Choice Problem. *Cerebral Cortex*, *25*(4), 972–982. https://doi.org/10.1093/cercor/bht286

Domenech, P., & Koechlin, E. (2015). Executive control and decision-making in the prefrontal cortex. *Current Opinion in Behavioral Sciences*, *1*, 101–106. https://doi.org/10.1016/j.cobeha.2014.10.007

Ferguson, T. S. (1989). Who Solved the Secretary Problem? *Statistical Science*, *4*(3), 282–289. https://doi.org/10.1214/ss/1177012493

Furl, N., & Averbeck, B. B. (2011). Parietal Cortex and Insula Relate to Evidence Seeking Relevant to Reward-Related Decisions. *Journal of Neuroscience*, *31*(48), 17572–17582. https://doi.org/10.1523/JNEUROSCI.4236-11.2011

Furl, N., Averbeck, B. B., & McKay, R. T. (2019). Looking for Mr(s) Right: Decision bias can prevent us from finding the most attractive face. *Cognitive psychology*, *111*, 1–14. https://doi.org/10.1016/j.cogpsych.2019.02.002

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, *61*(313), 35–73. https://doi.org/10.2307/2283044

Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding When to Decide: Time-Variant Sequential Sampling Models Explain the Emergence of Value-Based Decisions in the Human Brain. *Journal of Neuroscience*, *32*(31), 10686–10698. https://doi.org/10.1523/JNEUROSCI.0727-12.2012

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2017). Learning in the Repeated Secretary Problem. *arXiv preprint arXiv:1708.08831*, arxiv 1708.08831.

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning When to Stop Searching. *Management Science*, *66*(3), 1375–1394. https://doi.org/10.1287/mnsc.2018.3245

Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the Bayesians got their beliefs (and what those beliefs actually are): Comment on Bowers and Davis (2012). *Psychological Bulletin*, *138*(3), 415–422. https://doi.org/10.1037/a0026884

Guan, M., Lee, M. D., & Silva, A. (2014). Threshold Models of Human Decision Making on Optimal Stopping Problems in Different Environments. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*(36).

Guan, M., Lee, M. D., & Vandekerckhove, J. (2015). A Hierarchical Cognitive Threshold Model of Human Decision Making on Different Length Optimal Stopping Problems. *CogSci*.

Haber, S. N., & Knutson, B. (2010). The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology*, *35*(1), 4–26. https://doi.org/10.1038/npp.2009.129

Hill, T. P. (2009). Knowing When to Stop: How to gamble if you must—the mathematics of optimal stopping. *American Scientist*, *97*(2), 126–133.

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY, US, Farrar, Straus and Giroux.

Kiani, R., & Shadlen, M. N. (2009). Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science*, *324*(5928), 759–764. https://doi.org/10.1126/science.1169405

Knutson, B., Adams, C. M., Fong, G. W., & Hommer, D. (2001a). Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens. *Journal of Neuroscience*, *21*(16). https://doi.org/10.1523/JNEUROSCI.21-16-j0002.2001

Knutson, B., Fong, G. W., Adams, C. M., Varner, J. L., & Hommer, D. (2001b). Dissociation of reward anticipation and outcome with event-related fMRI. *Neuroreport*, *12*(17), 3683–3687.

Lee, M. D. (2006). A Hierarchical Bayesian Model of Human Decision-Making on an Optimal Stopping Problem. *Cognitive Science*, *30*(3), 1–26. https://doi.org/10.1207/s15516709cog0000_69

Lee, M. D., O'Connor, T. A., & Welsh, M. B. (2005). Decision-Making on the Full Information Secretary Problem. *Proceedings of the Twenty-Sixth Conference of the Cognitive Science Society*, 819–824.

Lin, Z., Nie, C., Zhang, Y., Chen, Y., & Yang, T. (2020). Evidence accumulation for value computation in the prefrontal cortex during decision making. *Proceedings of the National Academy of Sciences*, *117*(48), 30728–30737. https://doi.org/10.1073/pnas.2019077117

Ozer, D. J. (1990). Individual Differences and the Explanation of Behavior, In *Annals of Theoretical Psychology*. Boston, MA, Springer US.

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The Role of the Medial Frontal Cortex in Cognitive Control. *Science*, *306*(5695), 443–447. https://doi.org/10.1126/science.1100301

Ross, R., Pennycook, G., McKay, R., Gervais, W., Langdon, R., & Coltheart, M. (2016). Analytic cognitive style, not delusional ideation, predicts data gathering in a large beads task study. *Cognitive Neuropsychiatry*, *21*(4), 300–314. https://doi.org/10.1080/13546805.2016.1192025

Shu, S. B. (2008). Future-biased search: The quest for the ideal. *Journal of Behavioral Decision Making*, *21*(4), 352–377. https://doi.org/10.1002/bdm.593

Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior & Organization*, *35*(3), 309–332. https://doi.org/10.1016/S0167-2681(98)00051-1