METHODOLOGICAL REMARKS REGARDING OPTIMAL STOPPING TASKS AND THE IMPLICATIONS FOR SAMPLING BIASES

A PREPRINT

Didrika S. van de Wouw*

Department of Psychology
Royal Holloway, University of London
TW20 0EX, Egham, United Kingdom
sahira.vandewouw.2018@live.rhul.ac.uk

Ryan T. McKay

Department of Psychology
Royal Holloway, University of London
TW20 0EX, Egham, United Kingdom

Nicholas Furl

Department of Psychology
Royal Holloway, University of London
TW20 0EX, Egham, United Kingdom

September 14, 2021

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.

Keywords Optimal stopping · Decision-making · Bias · Bayesian

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).

*Corresponding author

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 ‘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 ($N_{\text{pilot1}} = 50$, $N_{\text{pilot2}} = 46$, $N_{\text{main}} = 294$). 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; $\mu_0 = -6.7402$, $\sigma_0^2 = 0.1038$). Log transformation was applied to the prices to approximate normality: a Shapiro-Wilk test of normality indicated that phone prices were not normally distributed ($p < .001$). 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 ($p < .001$). 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 Materials². 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 1). The reason why Figure 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.

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.

²<https://gorilla.sc/openmaterials/53623>

Figure 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.

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 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). 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/scarse 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 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 1).

Demographic information for participants enrolled into each of our six conditions in our Main Study can be found in Table 2. Because of a technical difficulty with the participant recruitment platform, we overshot our data collection

Table 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		✓				

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 1). We were surprised to find a null result in Pilot Study 1 (see Figure 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 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 2. However, despite the apparent differences in sampling bias between conditions (see Figure 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 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.

Table 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

Figure 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 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 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 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 ($p < .05$, Figure 2). In the *full* condition, which used Model 2, participants oversampled ($p < .01$, Figure 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.

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., 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*.
- 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.
- Prolific. (2014). Available at: <https://www.prolific.co>.
- RStudioTeam. (2020). RStudio: Integrated Development Environment for R.
- 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](https://doi.org/10.1016/S0167-4870(99)00038-0)