Published in final edited form as:

Cogn Psychol. 2019 June ; 111: 1–14. doi:10.1016/j.cogpsych.2019.02.002.

Looking for Mr(s) Right: Decision bias can prevent us from finding the most attractive face

Nicholas Furl^a, Bruno B. Averbeck^b, Ryan T. McKay^a

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

^bNIMH/NIH, 49 Convent Drive, MSC 4415, Bethesda, MD 20892-4415, United States

Abstract

In realistic and challenging decision contexts, people may show biases that prevent them from choosing their favored options. For example, astronomer Johannes Kepler famously interviewed several candidate fiancées sequentially, but was rejected when attempting to return to a previous candidate. Similarly, we examined human performance on searches for attractive faces through fixed-length sequences by adapting optimal stopping computational theory developed from behavioral ecology and economics. Although economics studies have repeatedly found that participants sample too few options before choosing the best-ranked number from a series, we instead found overlong searches with many sequences ending without choice. Participants employed irrationally high choice thresholds, compared to the more lax, realistic standards of a Bayesian ideal observer, which achieved better-ranked faces. We consider several computational accounts and find that participants most resemble a Bayesian model that decides based on altered attractiveness values. These values may produce starkly different biases in the facial attractiveness domain than in other decision domains.

Keywords

decision making; Bayesian modeling; facial attractiveness; mate choice; optimal stopping

1 INTRODUCTION

Many real-world decisions require optimal stopping. Accept a job offer or keep looking? Sell a stock now or wait for the price to rise? Buy a dress today or wait for a sale? In such “best choice” scenarios, agents must weigh the temptation to sample further options (so as potentially to improve on the current option) against the risk of missing the best option if too many are sampled. The classic illustration comes from astronomer Johannes Kepler’s search for a wife. After considering several candidates, Kepler returned to a previous candidate and

---

Corresponding author: Nicholas Furl, Department of Psychology, Royal Holloway, University of London, Egham, TW20 0EX, United Kingdom, nicholas.furl@rhul.ac.uk, +44(0)1784414390.

AUTHOR CONTRIBUTIONS. N.F., R.T.M. and B.B.A. designed studies. B.B.A. & N.F. implemented computational models. N.F. analyzed data. N.F. wrote the paper, and B.B.A. and R.T.M. provided critical revisions.

DECLARATIONS OF INTEREST: none

was duly rejected. Indeed, numerous investigators of this best-choice optimal stopping problem have associated it with mate choice (Eriksson & Strimling, 2009; Guan, Lee & Silva 2014; Todd, Billari & Simão, 2005; Todd & Miller, 1999), variously naming it the “fiancé(e)”, “marriage”, “dowry” or “fussy suitor” problem (Ferguson, 1989). Likewise, behavioral ecologists have, for decades, extensively studied how this decision structure relates to non-human animal mate choice using empirical studies (Valone et al., 1996) and theoretical models (Castellano et al., 2012; Collins, McNamara & Ramsey, 2006; Luttbegg, 1996; 2002; Janetos, 1980; Real, 1990). The computational treatments of this decision problem (e.g., Costa & Averbeck, 2015) typically address a version of this problem where prospects (e.g., potential partners) are limited in number (e.g., because of population size) and/or by search duration (e.g., an animal’s brief mating season). Agents facing this complication cannot simply set one aspiration threshold in advance and then wait for a sufficiently favorable option, as this strategy risks missing the highest-ranking available option (Kolling, Scholl, Chekroud, Trier & Rushworth, 2018). Optimal agents confronting fixed length sequences benefit by incorporating finite-choice horizons (Janetos, 1980).

Best-choice studies in humans are largely limited to fixed-length searches for best-ranked numbers in economic scenarios (e.g., find the car with the lowest mileage). These studies have established a pervasive and well-replicated finding: participants search too few options, compared to computational ideal observer (optimality) models (Bearden, Rapoport & Murphy, 2006; Costa & Averbeck, 2015; Seale & Rapoport, 1997; Seale & Rapoport, 2000; Sonnemans, 2000; Zwick, Rapoport, King Chung Lo & Muthukrishnan, 2003). The same finding arises in a closely related optimal stopping problem: the beads task and its variants. In the classic version of the beads task, participants infer the majority color of beads in a fictitious hidden jar before they have viewed an optimal number of samples of bead colors drawn from the jar (Furl & Averbeck, 2011; van der Leer, Hartig, Goldmanis, McKay, 2015; Hauser et al., 2017b). Vul and colleagues identified a number of other economic settings where participants make decisions based on undersampled probability distributions (Vul, Goodman, Griffiths & Tenenbaum, 2014). Mechanisms asserted to explain undersampling include robust heuristics (Todd & Miller, 1999), overweighting of evidence diagnosticity (van der Leer, Hartig, Goldmanis & McKay, 2017), excessive decision noise (Moutoussis et al., 2011), intrinsic search costs (Costa & Averbeck, 2015; Furl & Averbeck, 2011) and urgency signals (Hauser et al., 2017b).

Here, our main aim was to test whether humans also undersample in a more social decision scenario. We selected the mate choice domain, given that it motivated interest in best-choice problems in both mathematics (Ferguson, 1989) and behavioral ecology literatures (Castellano et al., 2012; Valone et al., 1996). Specifically, we focused on one important factor in human partner choices (out of many) – visual attractiveness. Although our initial expectation was that participants might undersample, there was some reason to suspect a different result might obtain. Some theories of animal sequential search choices (Janetos, 1980; Real, 1990), for example, assert that predispositions can bias choices toward phenotypes of a certain high quality (Beckers & Wagner, 2011; Ivy & Sakalu, 2007; Valone et al., 1996) and that these biases may be optimal on evolutionary scales (Cheng et al., 2014). Insofar as searches for facial attractiveness effectively trigger such dispositions,

biased preferences toward (relatively rare) high-quality partners might be expected to lengthen searches, instead of shorten them.

Our search task introduces a new approach to studying facial attractiveness choices. For the first time, computational models can be applied both as “ideal observer” optimality benchmarks and as mechanistic explanations of the computations human use when choosing attractive faces. The model we implement (Costa & Averbeck, 2015) combines prior information about possible option values with probabilistic learning to derive predictions of future outcome values. These predictions can then be used to compare currently-available option values against the probability of a better option appearing before the end of a fixed-length option sequence. This modeling arrangement is well-suited for “full information versions” of the best choice problem (Lee et al., 2006), like our facial attractiveness task. This model also benefits from being closely-related mathematically to the most commonly-used and well-established computational model of the beads task, (Averbeck, 2015; Furl & Averbeck, 2011; Hauser et al., 2017a; 2017b; Hauser, Moutoussis, Purg, Dayan & Dolan, 2018; Moutoussis, Bentall, El-Deredy & Dayan, 2011), as common computations underlie solutions to multiple optimal stopping problems. The model is also similar to several theoretical Bayesian models of non-human animal sequential choice (Castellano et al., 2012; Collins, McNamara & Ramsey, 2006; Luttbegg, 1996; 2002). The model draws from the same Bayesian framework often used to model behavioral and neural responses related to reward-guided decision making (Kolossa, Kopp & Fingscheidt, 2015; Solway & Botvinick, 2012). Likewise, Markov decision processes with dynamic thresholds implemented by this model are commonly used for decision models (Averbeck, 2015; Malhotra et al. 2018; Huang & Rao, 2013).

We report here three studies where we implemented a novel facial-attractiveness version of the best-choice decision task and compared human performance with that of a Bayesian ideal observer to measure bias. We tested which of multiple computational models best reproduced human behavior in two of these studies (that had sufficient data). Models included a biased values model as well as rival models that might produce similar behavior. Our primary interest was to establish whether the undersampling observed in economic domains indeed reflects a peculiarity of human probabilistic reasoning mechanisms that can infect any optimal stopping decision domain. If a different result obtains in our task, this calls into question whether biases in other new domains can be so easily predicted. We will also test gender differences as a secondary, more exploratory hypothesis, as some have argued that men and women make different mate choices (Fletcher, Kerr & Valentine, 2014).

2 MATERIALS AND METHODS

2.1 The three empirical studies

Informed consent was obtained from all participants in all studies, in accordance with the Declaration of Helsinki. Study 1 enrolled 49 participants, with the sample size based on our recent study on facial attractiveness choices (Furl, 2016). Participants chose a preferred sex (all 26 females and one male chose male faces, all others chose female faces). Because Study 1 offered roughly equally-sized face/participant sex groups, we used this dataset for an in-depth analysis of potential sex differences.

In phase 1, participants rated the attractiveness of 90 frontal, greyscale, youthful, neutral-expression faces (Burton, White & McNeill, 2010) of their chosen sex on a 9 point scale. Participants were asked to consider how much they would like to date the individuals in their ratings. Participants rated this image set three times. Participants' idiosyncratic preferences for each face were measured by the average of the three ratings from phase 1. We used these averages to rank faces in each phase 2 sequence and thereby assess decision performance. The use of personalized ratings to assess search performance protected these results against influences of extraneous variables, including individual participant and stimulus differences, that might affect the likelihood that a given face is chosen or not, apart from search strategy. Ratings also exposed participants and models to the prior distribution of attractiveness values that populated the sequences.

In phase 2, participants attempted to stop searching sequences of face images when they reached the most attractive face that they could. Participants were explicitly instructed to maximize the attractiveness of the faces in their choices. Depicted individuals were described as receptive potential partners so that participants would understand that they could not be rejected by any of their choices (the intention here was to avoid participants deliberately avoiding choosing the most attractiveness faces on the grounds that those individuals might be "out of the participant's league"). Five sequences of 12 faces each were organized using 60 faces pseudo-randomly selected from phase 1. Participants were informed that (a) not every face from phase 1 would be an option in phase 2, (b) they could not know the proportion of phase 1 faces used in phase 2, and (c) they could not know how many sequences there were and so any sequence might be the last chance to achieve an attractive date. The probability was small (0.12) that a specific phase 1 face would be sampled as an option in any given sequence and there was no guarantee that any face would appear in any sequence. The presentation of sequence options followed Costa & Averbeck (2015) as closely as possible, because this study had successfully replicated the classic undersampling effect using several economic, number-based scenarios and the same ideal observer model we used. The option screens in our studies (1) reminded participants of their number of remaining options in the current sequence and (2) showed, along the bottom of the screen, small "reminder" pictures of their refused options for that sequence. In Study 1, a new sequence was triggered upon choice of an option or if the last option was reached. This last option automatically became the chosen face for that sequence (once initially refused, options could never be returned to).

Power analysis of Study 1 data suggested that fewer than 20 participants would be sufficient for 95% power in Study 2. We enrolled 20 participants in this study. All 14 females and one male chose male faces, while all others chose female faces. Procedures were the same as Study 1, with the following modifications. Study 2 aimed to increase the amount of data per participant to facilitate our analysis and Bayesian model comparison of psychometric choice functions (28 phase 2 sequences cf. 5 in Study 1). The 426 rated faces (cf 90 in phase 1) rendered it even more improbable that a given phase 1 face would appear as a sequence option (only 8 options per sequence, hence a given phase 1 face had <2% chance of appearing in any sequence). Under these circumstances, participants should have plenty of experience with the distribution of attractiveness values. Waiting for any one specific face before stopping searching would be a highly irrational strategy. Study 2 also roughly equated

time spent on each sequence. After each choice, participants had to advance by keypress through grey squares that replaced the remaining pictures, so they could not finish sequences early by choosing an early option. In study 2, we also reinforced the reward value of choice using a feedback screen, displaying the participant's chosen face, the text "This is your date!", and a request to rate the reward value of the choice on a 9-point scale. As we needed many more faces, we sampled faces from a much larger set (Bainbridge, Isola & Oliva, 2013), choosing face images with happy expressions (which were numerous in this face set and allowed us to replicate with a different expression), which ranged in viewpoint degree between frontal and three-quarter view, and were color images of youthful individuals (apparently above 18 and less than 30, roughly approximating an undergraduate participant population) with circular grey masks.

Study 3 was originally designed to detect a between-participants effect of a mortality salience manipulation, $N=70$, based on a power analysis of a pilot study with $N=50$. The mortality salience group comparisons were not statistically significant (to be reported in separate manuscript). Nevertheless, the amount of data afforded by the large sample size was suitable for our goal of using choice data to compare participants with different theoretical models and so we applied this dataset to this purpose. Procedures were similar to Study 2, with the following exceptions. In phase 1, the 70 participants (60 female) rated the same 90 faces as in Study 1 (one female chose female faces, two males chose male faces, the rest chose opposite-sex faces) twice each and their averages were used to rank sequence options, as before. The two ratings were separated by either a mortality salience task or a dentistry imagination task (Rosenblatt, Greenberg, Solomon, Pyszcznski & Lyon, 1989). At phase 2, participants engaged with seven sequences of eight faces each.

2.2 The Bayesian modelling framework

Our model was customized for each participant. The model received as stimuli, and made decisions about, each corresponding participant's ratings for the faces in the same sequences as that participant experienced. The model, therefore, was also susceptible to general mate choice factors, which would have influenced that individual participant's attractiveness preferences, and used a search policy that sought to maximize these (participant-defined) factors by stopping searching at as highly-rated (by the participant) an option as possible. We chose a model that is closely-related mathematically to the model most commonly applied to a similar optimal stopping task, the beads task (Averbeck, 2015; Furl & Averbeck, 2011; Hauser et al., 2017a; 2017b; Hauser, Moutoussis, Purg, Dayan & Dolan, 2018; Moutoussis, Bentall, El-Deredy & Dayan, 2011). Our model also incorporates many elements previously considered for Bayesian models of animal sequential mate choice (Castellano et al., 2012; Collins, McNamara & Ramsey, 2006; Luttbeg, 1996; 2002).

Some previously-proposed optimality solutions, which define a "cutoff" based on an ideal search period (Dombrovsky & Perrin, 1994; Ferguson, 1989) have been mathematically proven to be optimal, given at least some of a restrictive set of assumptions that define the "secretary problem" version of the task. These assumptions include that (1) the agent cannot know or use information about the option value sampling distribution, but assumes that this distribution is stationary, (2) the agent knows only relative ranks of options but not their

absolute values and (3) the agent is rewarded only when choosing the highest-ranked sequence option. In contrast, our paradigm - a “full information problem” - has no need for these assumptions: (1) our participants generated the option sampling distribution themselves during phase 1; (2) our participants can directly perceive the absolute attractiveness value of each face, rather than its relative rank; (3) Participants are instructed to attempt to choose the most attractive face they could possible and had no way of knowing with certainty whether they actually achieved the highest-ranked face or not. Although it is plausible that cut-off heuristics are somewhat robust to violations of some of these assumptions (e.g., Bearden, 2006; Todd & Miller, 1999), there is no strong evidence that such heuristics are applicable to “full information problems” like our task. In contrast, our choice of model was specifically designed to provide normative results on best-choice tasks without making these restrictive assumptions (Costa & Averbeck, 2015).

Conceptually, the model we used computes values for the two possible actions (accept option versus decline/sample again) and acts on the higher-valued one. The action value for declining the current option can therefore be considered the current “aspiration threshold”, which the reward value of the current option must exceed for that option to be chosen. Because the action value of declining an option depends, in part, on probabilistic forecasts of future reward, the aspiration threshold is effectively dynamic and can change as the sequence progresses.

Mathematically, the model is based on a discrete time Markov decision process with continuous states. Action values combine reward values of potential options with internal representations of their probabilities, which are updated by every new sample of evidence. The utility u of each state s at time t across all available actions A_{S_t} is the maximum of the “action values” $Q(s_t, a)$, which depend on $r_t(s_t, a)$, the reward if action a is taken, and $p_t(j|s_t, a)$, the probability of transitioning to each state in the set of potential states S .

$$u_t(s_t, a) = \max_{a \in A_{S_t}} \left\{ r_t(s_t, a) + \int_S p_t(j|s_t, a) u_{t+1}(j) dj \right\}$$

Algorithmically, we use backward induction to compute utilities for each new state, because they depend on the utilities of subsequent states. Thus, we start by computing utilities for the final state s_N , which (because there are no state transition probabilities) is simply $u_N(s_N) = r_N(s_N)$ for all $s_N \in N$. Utilities for preceding states can then be computed as above, working backward from the last to the current one.

The model considers options as sampled from a Gaussian distribution with a $N\text{-Inv-}\chi^2$ prior (Gelman et al., 2004), which has four parameters: the prior mean μ_0 and variance σ_0^2 (set to the mean and variance of the attractive rating distribution in phase 1, which reflects the participants’ and model’s prior experience with the face set) and their respective degrees of freedom $\kappa_0 = 2$ and $\nu_0 = 1$ (set as in Costa & Averbeck, 2015). Each new sample yields a posterior distribution with new quantities $\mu_t, \sigma_t^2, \kappa_t, \nu_t$. The model’s probabilistic

representation of future outcome values is the distribution of state transition probabilities when the agent declines an option and chooses to sample another $p(j|s_t, a = \text{decline})$.

In addition to these representations of probabilities, action values also depend on representations of rewards. We defined a function R to map the outcome ranks onto reward values, whereby the model's reward was proportional to the corresponding participants' attractiveness rating. Using h as the relative rank of the current option (compared to declined options) and N as the number of outcome ranks in the sequence, we could compute the reward value of accepting the current option as

$$r_t(s_t, a = \text{accept}) = \sum_{i=1}^N p(\text{rank} = i) * R(i + (h - 1))$$

The corresponding reward value for declining an option is equivalent to the cost to sample, a quantity that in previous research (Furl & Averbeck, 2011; Costa & Averbeck, 2015) was usually zero (for an ideal observer model when no extrinsic cost was imposed by the experiment) or positive (when either the experiment costed samples or participants were assumed to experience an intrinsic aversion to sampling).

2.3 Theoretical model comparison

We compared participant data against the ideal observer model to test a null hypothesis that participants use a normative solution to this decision problem. The ideal observer implemented a cost to sample of zero, as there is no extrinsic cost in the study design. However, as reported in the Results, the human behavioral data suggested an oversampling bias, compared to the ideal observer. We therefore developed three candidate computational theories (described below), and tested how well they could predict the pattern of participants' choice thresholds. Models whose predictions conflict with basic patterns of the human behavioral data can be considered falsified (Palminteri, Wyart & Koechlin, 2017; Navarro, 2018), while models that can predict the participants' results continue to be viable candidates.

We determined participant choice thresholds by computing proportion choices in each of eight attractiveness bins and plotting them for every serial position. The resultant sigmoidal curves transition from zero choice (below-threshold attractiveness) to high levels of choice (above-threshold attractiveness). To estimate participant threshold locations for every serial position, we fitted logistic functions to these curves and computed points of subjective equality as the logistic inflection points. The inflection point estimates the attractiveness level where participants begin choosing faces. This analysis required considerable data to ensure there were sufficient data for stable proportion choice estimates for every one of these data points, for which Study 1 was insufficient (too few sequences per participant and too many sequence positions). We therefore performed our model comparisons by examining participant thresholds in Study 2 (which has many sequences per participant and fewer sequence positions) and 3 (we aggregated sequences over all 70 participants).

2.4. Theoretical model development

As reported in the Results, our analysis of psychometric functions showed that participants decreased their choice thresholds as sequences progressed (Lee, 2006). We therefore compared these thresholds against those of three rival computational theories that were designed to produce oversampling. Because the decreasing choice thresholds we observed in our studies cannot be produced by heuristics that do not employ probabilistic representations or use dynamic aspiration thresholds (e.g., Todd & Miller, 1999), we considered such heuristics *prima facie* falsified (Palmeri, Wyart & Koechlin, 2017) and we focus our more intensive model comparison on computational theories that rely on similar probabilistic mechanisms as the ideal observer model, as these use dynamic thresholds that in principle could reproduce participants' threshold setting behavior.

As described in the Introduction, we were inspired by theory from behavioral ecology suggesting that animals engaging in sequential mate searches may be predisposed to choose mates only with certain high-quality phenotypes (Janetos, 1980; Real, 1990; Valone et al., 1996; Cheng et al., 2014). We implemented a similar idea in the *biased values model*. Here, the attractiveness values of the sequence options are altered (relative to participants' original attractiveness ratings) and the model operates on these altered values. Such might be the case if participants normally use a normative probabilistic decision mechanism for decision making, but this mechanism is influenced by an external factor at the point of input. This external factor could be, among other things, a predisposition induced by the mate choice framing of the decision problem. To computationally instantiate such a state of affairs, we transformed the sequence values before they were submitted to the model using a logistic utility function, which was derived by fitting logistic functions (with maximum, sensitivity and bias as free parameters and minimum fixed to zero) to participants' averaged choice probabilities across the attractiveness bins (i.e., participant data in Fig. 3, averaged over serial positions). This transformation limits the influence that faces below a certain attractiveness level can have over choice, effectively raising attractiveness thresholds and leading to oversampling. However, the biased values model is not the only way a probabilistic Bayesian model like the one we consider might produce the oversampling observed in our participants. Even if the biased values model can predict participant behavior, other candidate models would need to be eliminated to draw a strong conclusion.

One such competitor model that we considered was a *sample reward model*. This model was motivated, in part, because a closely-related Bayesian model using a cost to sample term has been proposed to explain the classic undersampling effect both for number-based versions of the best-choice task (Costa & Averbeck, 2015) and for the beads task and its variants (Hauser et al., 2017b). This model characterizes the undersampling that occurs on these tasks as an intrinsic aversion to sample, or urgency signal (Furl & Averbeck, 2011; Hauser et al., 2017b). Using the same logic, we tested here whether searches might be extended because participants find viewing faces rewarding – akin to the apparently addictive qualities of on-line dating applications. We adapted our Bayesian model's cost-to-sample parameter (Costa & Averbeck, 2015) to implement an intrinsic reward value for sampling, which biases decisions in the direction of continued sampling. We employed a negative cost to sample

value (-0.035), which was selected to produce oversampling equal to that of the mean participant.

The third model we considered was the *attractive prior model*, which assumes that participants mis-represent the ratings distribution during phase 1, such that their prior belief is that faces will be, on average, more attractive than they were actually rated as being. Several factors might cause this. Participants' memory of the phase 1 set might be biased in favor of remembering more attractive faces than unattractive faces. Participants may rely on a different prior than the phase 1 distribution, perhaps acquired outside the study setting. Or participants may be susceptible to optimism bias, which brings an inflated prior expectation of positive future outcomes. Participants with the biased expectation that a highly attractive face might occur will wait longer to choose. We added a constant (0.5) to the prior mean μ_0 , to produce oversampling equal to that of the mean participant.

3 RESULTS

3.1 Optimal sampling? Comparisons with ideal observer

All three studies replicated the finding that participants sample more faces before choice than the ideal observer (Fig. 1a–c). Two-tailed t -tests, pairing participants with their corresponding models, showed highly significant effects (Study 1: $t(48) = 8.6$, $P < 0.001$; Studies 2 and 3: See Table 1). Despite searching less, the ideal observer achieved higher-ranked faces than participants (Fig. 1d–f). Two-tailed Friedman tests, pairing participants with their corresponding models, showed highly significant differences (Study 1: $\chi^2(1, N = 49) = 13.30$, $P < 0.001$; Studies 2 and 3: See Table 2).

Participants might oversample, as we observed, if they only chose rare, highly attractive faces. If so, participants should show a low, flat response rate across positions, with the highest proportion choices for the mandatory last image. This is because faces with the highest attractiveness values occur with the same low probability at every sequence position and, moreover, have a low probability of appearing anywhere in such short sequences. Fig. 2 confirms this response pattern. In contrast, the ideal observer chose most frequently about a third of the way through the sequence and thereby avoided the default last choice.

We also compared how attractiveness thresholds changed as sequences progressed. Proportion choices across attractiveness bins (Fig. 3) revealed sigmoidal curves, where choice can abruptly transition from zero choice to high levels of choice. The attractiveness value (bin) near this transition is taken as the choice threshold for that sequence position. We quantitatively estimated this threshold location (Fig. 4) as the point of subjective equality (i.e., inflection point of a fitted logistic function). Participants adopted nearly the highest possible attractiveness thresholds (i.e., bins 7, 8), consistent with a high-risk strategy of choosing rare, high-attractiveness faces. In contrast, the ideal observer used much lower choice thresholds. The ideal observer realistically estimates probabilities of attractiveness values so can accurately predict, for the remaining options, which attractiveness values are probable. The finite horizon that participants and the ideal observer adopted can be seen as a decline in thresholds toward sequence ends. Purpler curves in Fig. 3 are shifted leftward, compared to the bluer ones, and thresholds plotted in Fig. 4 are negatively-sloped.

3.2 Computational explanations for oversampling

To explain how participants deviated from the ideal observer model, we tested whether the three hypothetical Bayesian models described in Methods showed the same maladaptively high threshold pattern as the participants in Studies 2 and 3 (where we had sufficient data). When we compared participants, the ideal observer and these three new models using an omnibus analysis, we found highly significant main effects for both number of samples (ANOVA; Study 2: $F(3,54) = 80.34$, $P < 0.001$; Study 3: $F(3,204) = 56.93$, $P < 0.001$) and rank of chosen option (Friedman's test; Study 2: $\chi^2(3, N = 20) = 27.50$, $P < 0.001$; Study 3: $\chi^2(3, N = 70) = 27.45$, $P < 0.001$). *Post hoc* pairwise tests included both *t*-tests for number of samples and Friedman's tests for ranks – all two-tailed and Bonferroni-corrected for numbers of pairs within each study. In both studies (Tables 1 and 2), sample reward, attractive prior and biased values models resembled the participants in the sense that they oversampled, compared to the ideal observer. We had less evidence available to conclude any other pairwise differences between models or participants in the amount of sampling (Table 1). In contrast, sample reward and attractive prior models chose higher-ranked faces than participants, with less evidence for differences in chosen rank with the ideal observer model. The biased values model showed more ambiguous effects, with no detectable differences in rank of chosen faces with participants or other models (after Bonferroni correction), and a significant difference from the ideal observer only in Study 2.

These oversampling and rank measures verify that all three hypothetical models are viable explanations of participant oversampling behavior. However, these measures cannot easily distinguish whether any model explains participant behavior better than the others. We therefore also examined serial position effects. Fig. 2 shows that sample reward and attractive prior models had overlapping serial position curves, with slowly increasing choice rates as sequences progressed. In contrast, participants and the biased values model had overlapping serial position curves that maintained low choice rates throughout the sequence, resorting about 40% of the time to the last sequence option. When proportion choices was further broken down by serial position and attractiveness bin (Fig. 3), only the biased values model resembled the participants. Like participants, the biased values model used nearly the highest possible attractiveness thresholds (Figs. 3) and showed less of a threshold decline as sequences progressed, compared to the other models. This is also apparent in the points of subjective equality, where only the biased values model closely tracked the participants' thresholds (Fig. 4). Both sample reward and attractive prior models started with high thresholds, but then declined their thresholds more quickly than did the participants and biased values model.

We quantified this similarity of attractiveness thresholds for participants and biased values model by correlating participants' choices in all attractiveness bins and sequence positions (Fig. 3) with those of each model. These correlations are plotted in Fig. 4 and two-tailed pairwise Bonferroni-corrected significance tests are reported in Table 3. In Study 2, all three hypothetical models were better correlated with participants' pattern of choices than was the ideal observer. In both of the studies where these models were compared, the biased values model was better correlated with participants' pattern of choices than any of the other models.

3.3 Sex differences

As a secondary interest, we examined sex differences, given that sexes differ in attractiveness discriminability in contexts not requiring optimal stopping (Fletcher, Kerr & Valentine, 2014). Some theorists argue that men (but not women) should possess cognitive mechanisms adapted for minimising missed mating opportunities (Haselton & Bus, 2000; Haselton & Nettle, 2006; cf. McKay & Efferson, 2010; Perilloux & Kurzban, 2015). However, our data offer no strong evidence that oversampling bias varies by sex. Agent $\times$ face sex interactions were non-significant in all three studies ($P > 0.08$). In Study 1, where sexes were balanced, two-tailed t -tests using the participants who preferred opposite sex faces ($N = 46$, 52% female) showed no effect of sex on number of samples ($P = 0.10$) or rank ($P = 0.38$). Bayesian analysis, implemented in JASP (Wagenmakers et al., 2017), showed inconclusive Bayes Factors comparing a sex difference versus no sex difference model for sampling (0.882, error = 0.002) and rank (0.40, error = 0.022). We ranked both human and model choices according to each participant's own individual attractiveness ratings. As this already controls for individual differences in face preferences, sex differences in these preferences may also have been controlled.

4 DISCUSSION

4.1 Oversampling versus undersampling

Previous laboratory studies of optimal stopping have repeatedly shown that human participants sample fewer options than is computationally optimal. Participants in best-choice tasks are typically asked to choose high-ranked numbers in fictitious scenarios including buying a camera, renting an apartment, maximizing salary on a job offer and finding a credit card with a low interest rate (Costa & Averbeck, 2015; Seale & Rapoport, 1997; 2000; Sonnemans, 2000; Zwick et al., 2003). This best-choice task undersampling probably share a common computational mechanism with similar undersampling effects that have been replicated several times across the variants of a closely-related optimal stopping task, the beads task (Furl & Averbeck, 2011; Hauser et al., 2017b; van der Leer, Hartig, Goldmanis, McKay, 2015). In parallel, theoretical biologists have, for decades, considered sequential searches (analogous to best-choice tasks) to be one of the paradigmatic contexts for mate choice and therefore have proposed computational models similar to the one we used, which animals could use to solve sequential mate choice problems (e.g., Castellano et al., 2012). Inspired by this approach, we modified the classic human laboratory number-based best-choice task to involve instead a mate choice scenario with images of faces. Our results markedly departed from previous research: human participants sampled *more* faces than was optimal, rather than less. That is, participants sampled more and chose lower-ranked outcomes than the ideal observer. This was true, even though participants themselves generated the prior distribution of attractiveness values in phase 1, which should have indicated to them that extremely high-attractiveness faces would rarely appear in any short sequence. This main finding proved replicable across three studies that varied image sets, facial expressions, sequence lengths, numbers of sequences and other methods. The participants showed declining thresholds across serial positions, which accords with previous studies on searches for high-ranking numbers (Lee, 2006) and could arise from probabilistic representations with dynamic aspiration thresholds. Despite this threshold

decrease, participants kept thresholds too high throughout sequences. Consequently, they often refused high-ranked, but below-threshold, faces and continued sampling options until sequence termination.

Perhaps the most direct impact of this finding is to call into question existing theories of human decision making, which have been designed specifically to predict undersampling on optimal stopping tasks. Some theoretical proposals have applied the same types of computational models we consider here, incorporating a positive intrinsic cost to sample or urgency term to explain undersampling on the number-based best-choice task (Costa & Averbeck, 2015) and variants of the beads task (Furl & Averbeck, 2011; Hauser et al., 2017a). A closely-related explanation (Todd & Miller, 1999) supposes that participants will choose a sampling rate below the optimal one so long as it maintains near-optimal accuracy. Oversampling, however, would be surprising from the perspective of these theories, as they assume that lengthy searches are avoided because they consume resources and amplify risks (Furl & Averbeck, 2011; Todd & Miller, 1999). Nevertheless, we were motivated by similar reasoning to test the sample reward model, which generates oversampling from a sampling incentive, rather than a sampling cost. However, this model produced a different pattern of choice thresholds than the human participants, and so is not the most likely explanation for oversampling. In any case, there is hardly any existing consensus that favors this cost to sample explanation for undersampling. Alternative theoretical explanations for undersampling include overweighting of evidence diagnosticity (van der Leer, Hartig, Goldmanis & McKay, 2017) and excessive decision noise (Bearden, 2007; Moutoussis et al., 2011). Our data does not directly resolve controversy over the undersampling effect. However, our framework offers a new theoretical perspective. Our biased values model, which at least seems to explain oversampling, suggests that sampling biases can arise when option values are externally-weighted by other processes prior to their input into an otherwise-optimal probabilistic reasoning mechanism.

4.2 When does oversampling occur?

Why did this differential weighting and its consequent effect on choice thresholds and sampling rates, occur in our study but not previous ones? Many elements of our design, such as roughly normal option sampling distributions, numbers of options, numbers of sequences and reminders of previously rejected stimuli have all previously led to undersampling on the number-based task (Costa & Averbeck, 2015). Our three studies also replicate our finding across variations in experimental design elements such as numbers of sequences (from 5 to 28) and sequence options (8 or 15). Previous theories to explain the undersampling effect do not raise any predictions that simple design elements would negate the undersampling effect. Thus, any answer to the question of what causes oversampling would bring important theoretical implications. Our results here suggest that human sampling biases are not as predictable as previously believed and that further research will be needed to conclusively predict when participants will oversample, undersample or be optimal. Here, we discuss the two most obvious possibilities.

The most obvious difference between our paradigm and previous ones involves the mate choices that are implied when maximizing facial attractiveness. This possibility accords with

the biased values model, which was broadly inspired by biological theory proposing that mate-choosing animals set high thresholds, perhaps genetically determined (Cheng et al., 2014) on phenotypic variation when sequentially searching mates (Gibson, & Langren, 1996; Janetos, 1980; Real, 1990; Valone, Nordell, Giraldeau & Templeton, 1996), to the extent that some animals won't mate if sufficiently attractive options are not encountered. Our participants did not directly choose mates in our paradigm. They attempted to stop on the most attractive face possible, where attractive was explicitly defined as a desire to date the person. Moreover, mate choice does not uniquely occur only in the decision structure we presented to participants. Nevertheless, one possibility is that the activity of assessing attractiveness and the mention of a dating decision frame is sufficient to instigate mate choice predispositions. Such predispositions could then bias otherwise-optimal probabilistic choice mechanisms, in the way described by the biased valued model.

While the idea that the mate choice frame leads to oversampling via the biased values mechanism is a likely explanation, there are other differences between our paradigm and previous ones that require further study. A second obvious difference from previous work is the use of naturalistic image stimuli to convey option values to participants, instead of abstract stimuli like numbers conveying prices (Costa & Averbeck, 2015), relative ranks (Seale & Rapaport, 1997) or fictitious bead colors (Furl & Averbeck, 2011). A potential role for images in determining search strategy is intriguing because many real world searches (in addition to on-line dating applications) depend on natural images in general and face images in particular. In this case, hitherto unpredicted effects of decision domain may be more widespread than previously thought. Some real-world contexts, such as sequential eyewitness lineups or border control, require agents to assess sequentially-presented faces. These agents can only commit time and resources when sufficiently familiar faces are presented. Aside from face images, consumers engage with sequentially-presented naturalistic images when shopping for goods on-line and in catalogues. Many such situations involve choosing between committing to a pictorially depicted option and terminating search and forgoing an option to which it may be costly or impossible to return (e.g., a limited-time sale). Although one may have predicted only undersampling in these contexts previously, our results pose new empirical questions about whether different decision domains may induce different patterns of searching.

Are there also implications of our results for real-world mate searches in humans? Behavioral ecologists have long considered sequential search a fundamental mate choice context for animals (Beckers & Wagner, 2011; Castellano et al., 2012; Cheng et al., 2014; Collins, McNamara & Ramsey, 2006; Ivy & Sakalu, 2007; Luttbeg, 1996; 2002; Janetos, 1980; Real, 1990; Valone et al., 1996). Many scholars, starting with Kepler, assumed that human mate choice was a context with optimal stopping elements (Eriksson & Strimling, 2009; Guan, Lee & Silva 2014; Todd, Billari & Simão, 2005; Todd & Miller, 1999), giving rise to characterizations such as the “fiancé(e) problem” (Ferguson, 1989). Based on similar reasoning, previous research examined sequential mate choice contexts in paradigms simulating speed dating and on-line dating (Beckage, Todd, Penke, Asendorpf, 2009; Taubert et al., 2016) and has proposed that marriage rates are predictable based on best-choice optimal stopping logic (Todd et al., 2005). Our paradigm has enabled, for the first time, an application of computational probabilistic decision theory to facial attractiveness

choices. Despite this existing theoretical and empirical interest in mate choice as a sequential search problem, directly extending our results to human mate choices “in the wild” remains complicated and requires further data. The principal obstacle may be that there is not one “canonical mate choice context” but a vast diversity of mate-related contexts and decisions (e.g., commit to a date Saturday versus proposing marriage), not all of which are sequential. In general, our results may simply apply to any case where an agent must choose between committing exclusively to a receptive partner or moving on to explore other options. The frequency of these situations in the real world and whether they replicate the same decision biases remain open empirical questions.

4.3 Theories of decision making based on probabilistic representation

We tested rival computational theories to explain the oversampling bias. All three of our models successfully reproduced the participants’ oversampling bias and we distinguished among them based on their ability to predict participants’ dynamic aspiration thresholds across sequence positions. We considered probabilistic models with dynamic thresholds that could produce the participants’ threshold changes. Two of our models showed threshold changes that conflicted with those of the participants. Our sample reward model implemented an intrinsic reward for sampling more options, as might be expected if viewing attractive faces is rewarding. We also tested an “attractive prior” model in which the prior distribution had a maladaptively high mean value, as might be the case if the values of highly attractive individuals were especially well-encoded during the rating phase, if the prior distribution were skewed by experience outside the laboratory setting or if participants were subject to an optimism bias. While reward sample and attractive prior models produced some similar results to our participants, including oversampling biases, they also both manifested a precipitous threshold drop over sequence positions that did not match the more modest decline shown by our participants (Fig. 4). This occurs because thresholds in these two models can quickly re-adjust toward optimal as sequences progress. When there are fewer samples remaining to view, the sample reward model loses its prospect for future reward and incentive to continue sampling. Likewise, the attractive prior model learns from newly sampled values and thereby can quickly update and correct its probability distribution of values. Thus, both models start sequences with a bias but revert toward optimal as sequences progress. Although we have focused on some models that could plausibly reproduce oversampling and dynamic thresholds, our results cannot fully exclude all models or further modifications that might also reproduce our study data.

The evidence we have at hand favors a biased values model, in which the attractiveness values are non-linearly transformed prior to an otherwise-optimal decision process. The models we tested are, at best, approximations to the brain’s computations. Nevertheless, the core idea is that participants compute decision variables based on probabilistic representations of possible outcomes. This conclusion builds on evidence showing similar probabilistic reasoning mechanisms involved in other types of optimal stopping contexts (Castellano et al., 2012; Costa & Averbeck, 2015; Furl & Averbeck, 2011; Moutoussis et al., 2011). These mechanisms need not be specialized for solving best-choice tasks. Similar probabilistic representations could flexibly contribute to reward-guided decision-making

more generally (Averbeck, 2015; Gottlieb & Oudeyer, P-Y, 2018; Kolossa & Fingscheidt, 2015).

There are alternate theoretical approaches worth discussing. These were not included in our model comparison because they either (1) represent only partial theoretical accounts that don't fully specify computations participants hypothetically use to solve best-choice problems; (2) use static aspiration thresholds that cannot in principal reproduce the dynamic threshold strategies used by our participants or (3) they have not yet been fully developed as solutions for the "full information" version of the best-choice problem that we consider here. For example, cut-off heuristics propose that participants set a static aspiration threshold based on a learning period (Seale & Rapoport, 1997). These heuristics theories are based directly on a mathematical proof that provides an optimal solution to "secretary problems". Secretary problems are a special class of optimal stopping decisions and the optimal solutions mathematically holds only for its set of restrictive assumptions. Heuristic cut-off theories suppose that participants are aware of use this optimal solution and use a version of it that limits sampling slightly. This heuristic modification can lead to near-optimal performance in the context of secretary problem with less (potentially costly) sampling (Todd & Miller, 1999). However, the assumptions of the secretary problem do not hold in our paradigm (see Methods). Moreover, to accommodate our participant data, this heuristic would need further modification to explain why participants would increase rather than decrease sampling for our paradigm. More importantly, the heuristic would also need to be modified to have a dynamic aspiration threshold to replicate our participant's thresholds. Similarly, other models in the literature (Lee, 2006) would also need modification, as existing formulations also do not, at present, specify how participants compute their dynamic thresholds. On the whole, probabilistic representations like those we consider (Costa & Averbeck, 2015) are already well-suited by their design to full information problems like we consider. Further, they can provide *a priori* predictions about how participants compute their decisions that can reproduce the main features of our participants' data.

4.4 Conclusions

Our results suggest that surprisingly different biases might arise depending on decision context. Our evidence favoring the biased values model suggests that the attractiveness values (perceived phenotypes) that are entered into optimal probabilistic choice mechanisms are perhaps non-linearly related to those expressed in subjective ratings. Our approach aims to begin to disentangle influences of physical attractiveness and decision context in a way that allows us to bring to bear rigorous computational modelling methods that reveal new insights into human decision computations that could not otherwise be demonstrated. This approach continues the unification of a common cognitive framework that ties together theoretical development in mathematics, economics and behavioral ecology.

ACKNOWLEDGEMENTS.

We would like to thank Farah Atari, Laura Davies, Molly Haley, Iulia Nistor, Demi Riley-Carlaw, Galini Pantelidou and Alexandrina Vasilichi for assistance with data collection and Annabelle Gall for assistance developing stimuli.

"R.T.M. acknowledges funding support from the Cogito Foundation [grant number R10917]."

5. REFERENCES

- Averbeck BB. (2015). Theory of choice in bandit, information sampling and foraging tasks. *PLoS Computational Biology*, 11, e1004164. [PubMed: 25815510]
- Bainbridge WA, Isola P & Oliva A (2013). The intrinsic memorability of face images. *Journal of Experimental Psychology: General*, 142, 1323–1334. [PubMed: 24246059]
- Bearden JN (2006). A new secretary problem with rank-based selection and cardinal payoffs. *Journal of Mathematical Psychology*, 50, 58–59.
- Bearden JN (2007). On generalized secretary problems In *Uncertainty and risk* (eds Abdellaoui et al.), pp. 187–205. Springer.
- Bearden JN, Rapoport A & Murphy RO (2006). Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Science*, 52, 1437–1449.
- Beckers OM & Wagner WE Jr (2011). Mate sampling strategy in a field cricket: evidence for a fixed threshold strategy with last chance option. *Animal Behaviour*, 81, 519–527.
- Beckage N, Todd PM, Penke L & Asendorpf JB (2009). Testing sequential patterns in human mate choice using speed dating In *Proceedings of the 2009 Cognitive Science Conference* (eds. Taatgen N, van Rijn H), pp. 2365–2370. Cognitive Science Society.
- Burton AM, White D & McNeill A (2010). The Glasgow Face Matching Test. *Behavior Research Methods*, 42, 286–291. [PubMed: 20160307]
- Castellano S, Cadeddu G & Cermelli P (2012). Computational mate choice: Theory and empirical evidence. *Behavioural Processes*, 90, 261–277. [PubMed: 22410366]
- Cheng R, Seubert SM, & Wiegmann DD (2014). *Computational and Structural Biotechnology*, 10, 8–11.
- Collins EJ, McNamara JM & Ramsey DM (2006). Learning rules for optimal selection in a varying environment: Mate choice revisited. *Behavioral Ecology*, 17, 799–809.
- Costa VD & Averbeck BB (2015). Frontal-parietal and limbic-striatal activity underlies information sampling in the best choice problem. *Cerebral Cortex*, 25, 972–982. [PubMed: 24142842]
- Dombrovsky Y & Perrin N (1994). On adaptive search and optimal stopping in sequential mate choice. *The American Naturalist*, 144, 355–361.
- Eriksson K & Strimling P (2009). Partner heuristics in the lab: Stability of matchings under various preference structures. *Adaptive Behavior*, 17, 524–536.
- Ferguson TS (1989). Who solved the secretary problem? *Statistical Science*, 4, 282–289.
- Fletcher GJO, Kerr PSG, Li NP & Valentine KA (2014). Predicting romantic interest and decisions in the very early stages of mate selection: Standards, accuracy and sex differences. *Personality and Social Psychology Bulletin*, 40, 540–550. [PubMed: 24501043]
- Furl N (2016). Facial-Attractiveness Choices Are Predicted by Divisive Normalization. *Psychological Science*, 27, 1379–1387. [PubMed: 27565534]
- Furl N & Averbeck BB (2011). Parietal cortex and insula relate to evidence seeking relevant to reward-related decisions. *The Journal of Neuroscience*, 30, 17572–17582.
- Gelman A, Carlin JB, Stern HS & Rubin DB (2004). *Bayesian data analysis*. New York: Chapman and Hall/CRC.
- Gibson RM & Langen TA (1996) How do animals choose their mates? *Trends in Ecology and Evolution*, 11, 468–470. [PubMed: 21237923]
- Gottlieb J & Oudeyer P-Y (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*, 19, 7580770.
- Guan H, Lee MD, & Silva A (2014). Threshold models of human decision making on optimal stopping problems in different environments In Bello P, Guarini M, McShane M, & Scassellati B (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, pp. 553–558. Austin, TX: Cognitive Science Society
- Hahn AC, & Perrett DI (2014). Neural and behavioral responses to attractiveness in adult and infant faces. *Neuroscience & Biobehavioral Reviews*, 23, 591–603.
- Haselton MG & Buss DM (2000). Error management theory: A new perspective on biases in cross-sex mind reading. *The Journal of Personality and Social Psychology*, 78, 81–91. [PubMed: 10653507]

- Haselton MG & Nettle D (2006). The paranoid optimist: An integrative evolutionary model of cognitive biases. *Personality and Social Psychology Review*, 10, 47–66. [PubMed: 16430328]
- Hauser TU, Moutoussis M, Iannaccone R, Brem S, Walitza S, Drechsler R, Dayan P & Dolan RJ (2017a). Increased decision thresholds enhance information gathering performance in juvenile Obsessive-Compulsive Disorder (OCD). *PLoS Computational Biology*, 13, e1005440. [PubMed: 28403139]
- Hauser TU, Moutoussis M, NSPN Consortium, Dayan P & Dolan RJ (2017b). Increased decision thresholds trigger extended information gathering across the compulsivity spectrum. *Translational Psychiatry*, 7, 1296. [PubMed: 29249811]
- Hauser TU, Moutoussis M, Purg N, Dayan P & Dolan RJ (2018). BetaBlocker Propranolol Modulates Decision Urgency During Sequential Information Gathering. *The Journal of Neuroscience*, 38:7170–7178. [PubMed: 30006361]
- Ivy T & Sakaluk S (2007). Sequential mate choice in decorated crickets: females use a fixed internal threshold in pre- and postcopulatory choice. *Animal Behavior*, 74, 1065–1072.
- Janetos AC (1980). Strategies of Female Mate Choice: A Theoretical Analysis. *Behavioral Ecology and Sociobiology*, 7, 107–112.
- Kolling N, Scholl J, Chekroud A, Trier HA & Rushworth MFS (2018). Prospection, Perseverance, and Insight in Sequential Behavior. *Neuron*, 99, 10691082.
- Kolossa A, Kopp B & Fingscheidt T (2015). A computational analysis of the neural basis of Bayesian inference. *Neuroimage*, 106, 222–237. [PubMed: 25462794]
- Lee MD (2006). A hierarchical Bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, 30, 1–26.
- Luttbeg B (1996). A comparative Bayes tactic for mate assessment and choice. *Behavioral Ecology*, 7, 451–460.
- Luttbeg B (2002). Assessing the robustness and optimality of alternative decision rules with varying assumptions. *Animal Behaviour*, 63, 805–814.
- Malhotra G, Leslie DS, Ludwig CJH & Bogacz R (2018). Time-varying decision boundaries: Insights from optimality analysis. *Psychonomic Bulletin and Review*, 25, 971–996. [PubMed: 28730465]
- McKay R & Efferson C (2010). The subtleties of error management. *Evolution and Human Behaviour*, 31, 309–319.
- Moutoussis M, Bentall RP, El-Deredy W & Dayan P (2011). Bayesian modelling of jumping-to-conclusions bias in delusional patients. *Cognitive Neuropsychiatry*, 16, 422–447. [PubMed: 21480015]
- Navarro D (2018). Between the Devil and the Deep Blue Sea: Tensions Between Scientific Judgement and Statistical Model Selection. *Computational Brain & Behavior*, 1–7. [PubMed: 30906917]
- Palminteri S, Wyart V & Koechlin E (2017). The importance of falsification in computational cognitive modeling. *Trends in Cognitive Sciences*, 21, 425–433. [PubMed: 28476348]
- Perilloux C & Kurzban R (2015). Do men overperceive women's sexual interest? *Psychological Science*, 26, 70–77. [PubMed: 25413875]
- Real L (1990). Search theory and mate choice. I. Models of single-sex discrimination. *The American Naturalist*, 136, 376–404.
- Rosenblatt A, Greenberg J, Solomon S, Pyszczynski T & Lyon D (1989). Evidence For Terror Management Theory: I. The Effects of Mortality Salience on Reactions to Those Who Violate or Uphold Cultural Values. *Journal of Personality and Social Psychology*, 57, 681–690. [PubMed: 2795438]
- Seale D & Rapoport A (1997). Sequential Decision Making with Relative Ranks: An Experimental Investigation of the "Secretary Problem". *Organizational Behaviour and Human Decision Processes*, 69, 221–236.
- Seale D & Rapoport A (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Decision Making*, 13, 391–411.
- Sonnemans J 2000 Decisions and strategies in a sequential search experiment. *J. Econ. Psychol* 21, 91–102.

- Solway A & Botvinick M (2012). Goal-directed decision-making as probabilistic inference: A computational framework and potential neural correlates. *Psychological Review*, 119, 120–154. [PubMed: 22229491]
- Taubert J, Van der Burg E & Alais D (2016). Love at second sight: Sequential dependence of facial attractiveness in an on-line dating paradigm. *Scientific Reports*, 6, 22740. [PubMed: 26986828]
- Todd PM, Billari FC & Simão J (2005). Aggregate age-at-marriage patterns from individual mate-search heuristics. *Demography*, 42, 559–574. [PubMed: 16235613]
- Todd PM & Miller GF (1999). From pride and prejudice to persuasion: Satisficing in mate search In *Simple heuristics that make us smart* (eds. Gigerenzer G, Todd PM), pp 287–308. New York, NY: Oxford University Press.
- Valone TJ, Nordell SE, Giraldeau L-A & Templeton JJ (1996). The empirical question of thresholds and mechanisms of mate choice. *Evolutionary Ecology*, 10, 447–455.
- van der Leer L, Hartig B, Goldmanis M & McKay R (2015). Delusion proneness and ‘jumping to conclusions’: relative and absolute effects. *Psychological Medicine*, 45, 1253–1262.
- van der Leer L, Hartig B, Goldmanis M & McKay R (2017). Why do delusion-prone individuals “jump to conclusions”? An investigation using a non-serial data gathering paradigm. *Clinical Psychological Science*, 5, 718–725.
- Vul E, Goodman N, Griffiths T & Tenenbaum J (2014). One and Done: Optimal decisions from very few samples. *Cognitive Science*, 38, 599–637. [PubMed: 24467492]
- Wagenmakers EJ, Love J, Marsman M, Jamil T, Ly A, Verhagen J, ... Morey RD (2017). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic Bulletin and Review* doi: 10.3758/s13423-017-1323-7. [Epub ahead of print].
- Huang Y & Rao RPN Reward optimization in the primate brain: A probabilistic model of decision making under uncertainty. *PLoS One*, 8, e53344. [PubMed: 23349707]
- Zwack R, Rapoport A, Chung, Lo A, Muthukrishnan AV (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, 22, 503–519.

Figure 1.

Oversampling biases. Mean number of faces that participants and models sampled until choice are shown for Studies 1 (a), 2 (b) and 3 (c). Mean attractiveness rank in each sequence of chosen faces is shown for Studies 1 (d), 2 (e) and 3 (f). Distributions are indicated by kernel densities overlaid with individual data points with white horizontal lines denoting mean values.

Figure 2.

Serial position effects. Mean proportion choices for each sequence position are shown for participants and models. Shaded areas show 95% confidence intervals. Panels a-c are Studies 1-3.

Figure 3. Proportion choices for each attractiveness bin plotted separately for each serial position.

Figure 4.

Model comparisons. Points of subjective equality (attractiveness thresholds, measured as inflection points of logistic functions fitted to data in Fig. 3) are shown for Studies 2 (a) and 3 (b). To directly compare the models' ability to explain participant behavior, correlations between participants' behavior and that of each model were computed for the patterns of data shown in Fig. 3. For Study 2, correlations were computed between data from each model and performance from its corresponding participant. Shown are average correlations and their 95% confidence intervals over these participant/model pairs (c). In Study 3, each correlation was computed between data from a model and performance aggregated over all participants.

Table 1.

Pairwise two-tailed *t*-tests for differences between participants and models for mean number of samples until choice.

		participants	sample reward	attractive prior	biased values
sample reward	study 2	$d = .02$ $t(19) = .1$ $P = 1$			
	study 3	$d = .07$ $t(69) = -4.0$ $P = 1$			
attractive prior	study 2	$d = .02$ $t(19) = .1$ $P = 1$	$d = <.001$ $t(19) < .1$ $P = 1$		
	study 3	$d = .02$ $t(69) = .1$ $P = 1$	$d = .09$ $t(69) = -.7$ $P = 1$		
biased values	study 2	$d = .53$ $t(19) = -2.3$ $P = 0.34$	$d = .66$ $t(19) = -2.9$ $P = 0.1$	$d = .66$ $t(19) = -2.8$ $P = 0.1$	
	study 3	$d = .14$ $t(69) = -1.2$ $P = 1$	$d = .22$ $t(69) = -1.8$ $P = .72$	$d = .13$ $t(69) = -1.1$ $P = 1$	
ideal observer	study 2	$d = 2.97$ $t(19) = 13$ $P < .001$	$d = 3.76$ $t(19) = 16.4$ $P < .001$	$d = 3.1$ $t(19) = 13.5$ $P < .001$	$d = 3.25$ $t(19) = 14.1$ $P < .001$
	study 3	$d = 1.54$ $t(69) = 12.8$ $P < .001$	$d = 2.17$ $t(69) = 18$ $P < .001$	$d = 2.01$ $t(69) = 16.7$ $P < .001$	$d = 1.87$ $t(69) = 15.5$ $P < .001$

Note: *P*-values are Bonferroni-corrected. Tests printed in bold are significant after Bonferroni correction for number of pairs in each study. Tests presented in normal typeface are significant only when uncorrected. Tests presented in gray are non-significant, with or without correction.

Table 2.

Pairwise two-tailed Friedman tests for differences between participants and models for mean rank in sequence of chosen option.

		participants	sample reward	attractive prior	biased values
sample reward	study 2	$d = 1.27$ $\chi^2 = 16.2$ $P < .001$			
	study 3	$d = .48$ $\chi^2 = 10.6$ $P = .001$			
attractive prior	study 2	$d = 1.4$ $\chi^2 = 12.8$ $P = .003$	$d = .08$ $\chi^2 < .1$ $P = 1$		
	study 3	$d = 0.54$ $\chi^2 = 11.5$ $P < .001$	$d = .11$ $\chi^2 = 5.9$ $P = .15$		
biased values	study 2	$d = .09$ $\chi^2 = .5$ $P = 1$	$d = .91$ $\chi^2 = 5$ $P = 0.2$	$d = 1.09$ $\chi^2 = 5$ $P = .25$	
	study 3	$d = .34$ $\chi^2 = 6.1$ $P = .13$	$d = .17$ $\chi^2 = .1$ $P = 1$	$d = .21$ $\chi^2 = 2.6$ $P = 1$	
ideal observer	study 2	$d = 1.44$ $\chi^2 = 12.8$ $P = .003$	$d = .45$ $\chi^2 < 7.2$ $P = .07$	$d = .36$ $\chi^2 = .8$ $P = 1$	$d = 1.05$ $\chi^2 < 0.1$ $P < .001$
	study 3	$d = 0.82$ $\chi^2 = 22.9$ $P < .001$	$d = .39$ $\chi^2 = 5.4$ $P = .2$	$d = .27$ $\chi^2 = .4$ $P = .53$	$d = .46$ $\chi^2 = 4.4$ $P = .35$

Note: P -values are Bonferroni-corrected. Tests printed in bold are significant after Bonferroni correction for number of pairs in each study. Tests presented in normal typeface are significant only when uncorrected. Tests presented in gray are non-significant, with or without correction.

Table 3.

Pairwise model comparisons between models and participants of correlation coefficients of functions relating sequence position and proportion choices (Fig. 3)

		sample reward	attractive prior	biased values
attractive prior	study 2	$d = .27$ $t(19) = -1.2$ $P = 1$		
	study 3	$z = -8$ $P = 1$		
biased values	study 2	$d = 1.02$ $t(19) = -4.5$ $P = .002$	$d = .9$ $t(19) = -3.9$ $P = .006$	
	study 3	$z = -3.5$ $P = .003$	$z = -2.7$ $P = .04$	
ideal observer	study 2	$d = .96$ $t(19) = 4.2$ $P = .003$	$d = 1.15$ $t(19) = 5.0$ $P < .001$	$d = 1.77$ $t(19) = 7.7$ $P < .001$
	study 3	$z = 1.1$ $P = 1$	$z = 1.9$ $P = .36$	$z = 4.6$ $P < .001$

Note: P -values are Bonferroni-corrected. Tests printed in bold are significant after Bonferroni correction for number of pairs in each study. Tests presented in gray are non-significant, with or without correction.