



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Perceiving prospects properly

Citation for published version:

Steiner, J & Stewart, C 2014 'Perceiving prospects properly' ESE Discussion Papers, no. 245, Edinburgh School of Economics Discussion Paper Series.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Publisher Rights Statement:

© Steiner, J., & Stewart, C. (2014). Perceiving Prospects Properly. (Edinburgh School of Economics Discussion Paper Series).

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.





Edinburgh School of Economics
Discussion Paper Series
Number 245

Perceiving Prospects Properly

Jakub Steiner

(CERGE-EI and University of Edinburgh)

Colin Stewart

(University of Toronto)

Date

September 2014

Published by

School of Economics
University of Edinburgh
30 -31 Buccleuch Place
Edinburgh EH8 9JT
+44 (0)131 650 8361
<http://edin.ac/16ja6A6>



THE UNIVERSITY *of* EDINBURGH

Perceiving Prospects Properly*

Jakub Steiner[†]

CERGE-EI and University of Edinburgh

Colin Stewart[‡]

University of Toronto

August 22, 2014

Abstract

When an agent chooses between prospects, noise in information processing generates an effect akin to the winner’s curse. Statistically unbiased perception systematically overvalues the chosen action because it fails to account for the possibility that noise is responsible for making the preferred action appear to be optimal. The optimal perception patterns share key features with prospect theory, namely, overweighting of small probability events (and corresponding underweighting of high probability events), status quo bias, and reference-dependent S-shaped valuations. These biases arise to correct for the winner’s curse effect.

1 Introduction

There is considerable evidence that human perception of reality is noisy and biased.¹ While randomness can be understood as a technological limitation of human cognition, systematic behavioral biases, such as those documented in the psychological experiments of Kahneman and Tversky (1979), are more puzzling. Since there is no obvious reason why natural or

*We thank Michal Bauer, Andrew Clausen, Olivier Compte, Ed Hopkins, Tatiana Kornienko, David Levine, Filip Matějka, Fabio Michelluci, Nick Netzer, Motty Perry, Ariel Rubinstein, József Sákovics, Balázs Szentes, Tymon Tatur, and participants in various seminars and conferences for their comments. Maxim Goryunov, Ludmila Matysková, Jan Šípek, and Regina Tikhbatullina provided excellent research assistance.

[†]email: jakub.steiner@cerge-ei.cz

[‡]email: colinstewart@gmail.com

¹McFadden (1999) summarizes the experimental evidence as follows: “humans fail to retrieve and process information consistently. . . . These failures may be fundamental, the result of the way human memory is wired. I conclude that perception-rationality fails, and that the failures are systematic, persistent, pervasive, and large in magnitude.”

cultural evolution could not remove these biases, their prevalence suggests that they serve a purpose.

This paper argues that perception biases arise as second-best solutions when some noise in information processing is unavoidable. In particular, we show that overweighting of small probability events optimally mitigates errors due to randomness. Other biases captured by prospect theory—namely, status quo bias, and reversal of risk preferences across gains and losses—serve the same purpose in related settings.

In addition to providing an explanation for commonly observed cognitive biases, our analysis offers predictions about how the biases vary. Our model also provides a framework for conceptualizing errors in decision-making, allowing us to consider, for example, whether overweighting of small probabilities is a mistake or an optimal heuristic. Finally, our results demonstrate how explicitly modelling the structure of decision-making can illuminate patterns of observed behavior.²

Our model separates decision-making into two stages. At the first stage, the decision-maker observes the parameters of the decision problem and encodes them using a perception strategy. The encoded values are then subject to stochastic noise. At the second stage, the decision-maker chooses an action based on these noisy values of the parameters. The noise can be interpreted as physiological randomness in the functioning of the brain, as failing to remember or keep track of all relevant information during decision-making, or as random computational errors. Each of these cases can be viewed as a loss of information during the decision process. Our main focus is on the optimal design of the perception strategy: given that noise will prevent the agent from using the true values in the second stage, how should those values be encoded beforehand?³

One natural perception strategy to consider is the unbiased one that gives rise, on average, to the correct parameter values after the noise is added. We argue that the unbiased strategy suffers from a problem akin to the “winner’s curse” that makes it suboptimal. Just as a bidder in a common value auction should condition her value on winning, the design of

²Our approach lies between two opposing positions on the relevance of cognitive processes in economic theory. On the one hand, as advocated by Gul and Pesendorfer (2008), we model and rationalize putative behavioral mistakes as optimal strategies under informational constraints. On the other hand, we agree with Camerer et al. (2005) that explicit modeling of the cognition process may enrich rational choice theory and inform welfare debates.

³We view the perception strategy as being applied subconsciously and optimized through evolution rather than through conscious reasoning. Kirkpatrick and Epstein (1992) present experimental evidence suggesting that subconscious distortions drive choice even when subjects correctly identify objective probabilities. See also Camerer et al. (2005) for a discussion of conscious and subconscious processing of probabilities.

the perception strategy should condition on which action is chosen. Unbiased perception fails to account for the possibility that noise is responsible for making the preferred action appear to be optimal. Biases in perception can correct for this winner’s curse by generating a more cautious evaluation of actions.⁴

The intuition for our results is simplest in the case of the status quo bias. Consider an agent who chooses between the status quo and an alternative action. The agent’s perception is chosen from a class of strategies differing only in the degree of status quo bias, i.e., in the extent to which the perception of the status quo reward is exaggerated. In particular, the strategy with no status quo bias yields unbiased perception of rewards, whereas with a nonzero bias, the agent’s average perception systematically favors one of the two actions. The perception problem consists of choosing the degree of bias that maximizes the expected reward the agent receives across all possible realizations of the binary decision problem. It turns out that the unbiased strategy is (generically) suboptimal: the optimal perception strategy is unbiased *conditional on the two options being perceived as equally attractive*, which implies that, *unconditionally*, it is biased.

Suppose that the average status quo is better than typical rewards from the alternative action (as one might expect if the status quo results from previous optimizing choices). Then unbiased perception leads to a winner’s curse because, conditional on perceiving the alternative as optimal, the agent overvalues it. As a result, the optimal status quo bias is positive, correcting for the winner’s curse; optimal perception makes the agent cautious about the alternative because it may only appear to be optimal due to an error.

The main focus of this paper is the perception of probabilities. As in the preceding example, unbiased perception of probabilities leads to a winner’s curse since errors that increase the relative attractiveness of an action make that action more likely to be chosen. We argue that overweighting small probabilities (and underweighting large ones) mitigates the overoptimism stemming from the winner’s curse.

Probability distortions may, at first blush, seem unlikely to help. Biasing probabilities does not, on average, make the agent more pessimistic or optimistic; exaggeration of small probabilities makes the agent less inclined to fly for fear of an accident, but more inclined to play casino games that offer a small probability of a large reward. However, such a bias does tend to make the agent more pessimistic about *attractive* lotteries. The reason is that lotteries perceived as valuable are much less likely to share the structure of a casino game

⁴See Compte and Postlewaite (2012) for a similar interpretation in a different setting.

than that of a decision to fly. Compare two lotteries offering the same expected reward: one a “flight lottery” that gives a high probability gain and a low probability loss, and the other a “casino lottery” that gives a low probability gain and a high probability loss. For any given loss, the two lotteries can have the same expected reward only if the low probability of a gain in the casino lottery is compensated with a very high reward. If very high rewards are rare, then attractive lotteries are typically like the flight lottery. Exaggerating low probabilities therefore tends to increase the weight given to losses, reducing the perceived value of the most attractive lotteries.

Is overweighting of small probabilities (and underweighting of large ones) a mistake? On the one hand, since such a bias perception is an optimal response to subsequent information loss, the agent would be worse off on average if he “debiased” his perception across all decision problems. On the other hand, a globally optimal perception strategy may perform poorly in some decision problems. In particular, the ex ante optimal perception strategy performs badly when the agent faces the casino lottery described above. Since the casino lottery is unlikely to be an attractive option, the ex ante optimal strategy introduces relatively large perception errors in lotteries of that form. Ex post, an outside observer who knows that the agent faces such a lottery could reasonably characterize the agent’s perception bias as a mistake because it is suboptimal given the observer’s information. Section 6 discusses these issues in more detail.

Our model offers predictions about the impact of the status quo on perception biases. We find that the optimal status quo bias switches from positive to negative if the average value of the status quo becomes lower than that of rewards from alternative actions. In contrast, the optimal bias in probability perception has the same inverted S-shape for high and low status quo values. This suggests that the direction of the status quo bias should vary consistently with an agent’s well-being, while the broad shape of probability weighting should not.

Although our model of the cognitive process is stylized, the neuroscientific literature offers some support for a two-stage, noisy decision process. Glimcher (2009) describes the emerging neuroeconomic consensus that the choice system in primates “involves a two-stage mechanism. The first of these stages is concerned with the valuation of all goods and actions; the second is concerned with choosing. . . [from] the choice set.” Tobler et al. (2008) document that probabilities are simultaneously encoded in more than one area of the brain, and that neuronal coding of probabilities in areas associated with probabilistic decision-making shows an inverted S-shaped pattern. Bossaerts et al. (2009) discuss evidence that

there are at least two imperfectly correlated brain signals involved in the choice process, one for assessing value, the other for the choice itself. More broadly, Glimcher (2005) surveys a body of evidence suggesting fundamental randomness in the activity of the brain.

Our paper fits into the literature on the principal-agent approach to evolution (see, e.g., Robson, 2001b; Samuelson and Swinkels, 2006; Robson and Samuelson, 2011). Robson (2001a), Rayo and Becker (2007), and Netzer (2009) study the evolutionary design of incentives for agents who cannot process information perfectly. They find that the optimal incentives are steeper at ranges of stimuli that the agent encounters more frequently, which can be interpreted as allocating greater attention to more common problems.⁵ Our results can also be understood in terms of optimal attention allocation, but extended to choice under uncertainty and using a different model of information processing.

Several papers study foundations for the biases captured in prospect theory. Herold and Netzer (2010) argue that inverted S-shaped probability weighting is an optimal response to S-shaped valuation of rewards. Similarly, Frenkel et al. (2012) view the endowment effect as a heuristic benefitting agents who suffer from the winner’s curse in bilateral trade. In contrast with those two papers, we derive optimal distortions in perception in the absence of frictions in other dimensions of the decision process. Woodford (2012a,b) studies optimal perception using insights from the rational inattention literature. Woodford’s analysis focuses on a relatively simple objective (namely, minimization of the mean square error) while allowing for a rich class of perception strategies. In contrast, we focus on maximization of expected rewards, and identify systematic deviations relative to the mean-square-error-minimizing perception.

Compte and Postlewaite (2012) study optimal heuristics for choice under uncertainty, and identify conditions under which a decision-maker exhibits “cautiousness” toward less certain outcomes. While our results can be interpreted similarly as a form of cautiousness, and our status quo bias example overlaps with their model, we differ significantly in terms of focus and modelling approach. In particular, in their model, cautiousness is optimal in a class of relative simple strategies, whereas in ours it is a feature of the best response to imperfections in information processing.

⁵Friedman (1989) provides an early analysis of the attention allocation problem using a reduced form of evolutionary optimization.

2 Model

An agent faces a series of binary decision problems, one in each period $t = 0, 1, \dots$. In each of these problems he chooses between a *status quo* that delivers payoff s , and a lottery that pays a reward r_1 or r_2 with respective probabilities p and $1 - p$. While all of these parameters are observable to the agent, he may make suboptimal choices due to errors in information processing. We distinguish between two stages of decision-making. In the first stage, an observation center observes the probability p and sends a message $m(p) \in [\underline{m}, \overline{m}]$ to a decision center. The message is subject to random noise captured by a term ε drawn from an interval; we view ε as resulting from physical noise within the agent's brain, from a failure to retain information, or from computational errors. The message m and noise ε combine to form the *perceived probability* (or simply the *perception*) $q = c(m, \varepsilon) \in [0, 1]$. The function c captures both the physical properties of the communication channel and the way in which the decision center decodes the arriving stimulus. In the second stage, the decision center chooses the lottery if $qr_1 + (1 - q)r_2 > s$, and chooses the status quo otherwise. Thus, for any lottery $\ell = (r_1, p; r_2, 1 - p)$ and perception q , the agent receives expected payoff

$$f(\ell, q) = \begin{cases} pr_1 + (1 - p)r_2 & \text{if } qr_1 + (1 - q)r_2 > s, \\ s & \text{otherwise.} \end{cases}$$

Figure 1 summarizes the decision process.

The values of the rewards and the status quo are measured in terms of utilities that represent expected fitness, and incorporate risk preferences. Although the function c and the choice rule of the decision center are fixed, we implicitly allow for the possibility that behavior is also optimized at the decision stage since, given any perception strategy, the optimal decision rule corresponds to one of the functions c that we consider (see Section 4.1 for details). Moreover, by not insisting on optimality of c , we allow for—but do not require—the existence of constraints in the evolution of the decision rule.

Nature chooses a *perception strategy* $m(p; s)$ to maximize the agent's ex ante expected payoff for a given distribution over lotteries.⁶ Draws of p and (r_1, r_2) are independent, and independent of the noise ε . In addition, the distribution of (r_1, r_2) is symmetric and the density ψ of p is symmetric around $1/2$, in accordance with the idea that the indices have

⁶When it is not needed, we often drop the argument s from the notation.




Figure 1: The two-stage decision process with interim noise.

no intrinsic meaning. An optimal perception strategy $m^*(p; s)$ satisfies

$$m^*(\cdot; s) \in \arg \max_{m(\cdot)} E[f(\ell, c(m(p), \varepsilon))], \quad (1)$$

where the expectation is over the noise ε and the lottery ℓ . An optimal strategy always exists. If there are multiple optimal strategies then all of our results hold for each such strategy. We therefore ignore potential multiplicity and often refer to “the” optimal strategy.

2.1 The value of the status quo

The optimal perception strategy depends on the status quo s . Although we take s to be exogenous in our static model, one can view it as resulting from previous optimizing choices by the agent. We argue that this makes the status quo likely to exceed the average rewards available to the agent. This section offers a simple dynamic model formalizing this argument.

The lottery available to the agent is i.i.d. across periods. The agent forms a perception $q_t = c(m^*(p_t; s_t), \varepsilon_t)$ according to the optimal strategy m^* . If, based on this perception, the agent chooses the status quo, he enters the next period with status quo $s_{t+1} = s_t$. If he chooses the lottery, then the new status quo s_{t+1} is equal to r_{1t} with probability p_t and equal to r_{2t} with probability $1 - p_t$. To summarize,

$$s_{t+1} = \begin{cases} s_t & \text{if the agent chooses } s_t, \\ r_{1t} & \text{if the agent chooses } \ell_t \text{ and } r_{1t} \text{ is drawn,} \\ r_{2t} & \text{if the agent chooses } \ell_t \text{ and } r_{2t} \text{ is drawn.} \end{cases}$$

For any given initial status quo s_0 , we characterize the optimal perception in the long-run,

as t grows large.

Although the agent chooses myopically in each period, one can interpret the values r_{1t} , r_{2t} , and s_t as continuation values in a dynamic optimization problem, in which case the choice rule optimizes over his lifetime. The only complication is that continuation values are endogenously determined by the agent's strategy. At an optimum, however, our results hold under either interpretation.

Under a mild condition on the perception strategy m^* (which we verify below), the status quo eventually becomes large relative to the distribution of rewards in the lottery. Intuitively, any reasonable perception strategy generates an upward drift in the status quo because better prospects are more likely to be chosen, which in turn tends to make the status quo grow large. Because of this result, to identify optimal behavior in the long-run, it suffices to focus on decision problems in which the status quo is high.

Proposition 1. *Suppose that the perception $q = c(m^*(p; s), \varepsilon)$ is nondecreasing in p . Then for each $s^* \in \mathbb{R}$,*

$$\lim_{t \rightarrow \infty} \Pr(s_t > s^*) = 1.$$

Proofs are in the appendix.

3 Special case

Before analyzing the general model, in this section we illustrate the main result in a relatively simple special case with a particular distribution of rewards and additively separable noise. We relax many of the following assumptions in Section 4.

In each period, the rewards r_1 and r_2 are independently drawn from the standard normal distribution. For each message m , the perception is given by $q = m + \varepsilon$, where the noise ε attains values σ and $-\sigma$, each with probability $1/2$, and $\sigma \in (0, 1/2)$. To avoid complications due to boundary effects, the density $\psi(p)$ has support contained in $[\sigma, 1 - \sigma]$, and the message space is $[\sigma, 1 - \sigma]$, ensuring that the perception q is always in $[0, 1]$.

In the absence of noise, the perception optimization problem is trivial: the unbiased perception strategy $m(p) \equiv p$ achieves the first-best. When there is noise, the optimal perception strategy exhibits systematic biases.

Theorem 1. *The optimal perception strategy $m^*(p; s)$ is nondecreasing in p . Furthermore, if $s > 3^{1/4}$, then the agent overstates small probabilities and understates large probabilities; that is, for all $p \in (\sigma, 1 - \sigma) \setminus \{1/2\}$, $|m^*(p; s) - 1/2| < |p - 1/2|$.*




Figure 2: The optimal perception strategy $m^*(p; s)$ (solid curve) for status quo $s = 2$ relative to the unbiased strategy $m(p; s) \equiv p$ (dashed line).

Figure 2 depicts the optimal probability perception for a particular value of the status quo.

Although the strategy $m(p)$ describes internal communication within the agent, the perception is in principle observable in an experiment. By varying the rewards, an experimenter can recover the subject’s stochastic probability perception $q = m(p) + \varepsilon$ of the objective probability p . The average perception $E[q | p]$ across many repetitions of the experiment is equal to $m(p)$. Theorem 1 therefore indicates that an agent using the optimal strategy m^* will be seen to be overweighting small probabilities and underweighting large ones.


Note that Nature does not condition the perception of p on the rewards in the lottery. If the perception could depend on rewards, the first-best could be achieved by effectively making the observation center compute the optimal action and then send an extreme message to the decision center indicating which action to take. By requiring the perception strategy $m(p; s)$ to depend only on p and s , we constrain Nature to choose a heuristic that performs well on average across all possible rewards. This approach is consistent with neuroscientific evidence that probabilities and rewards are processed separately (see Berns and Bell, 2012).

3.1 Intuition

This subsection provides an intuitive explanation for why the optimal perception strategy overstates low and understates high probabilities.

When does a small change in perception affect choice? If the expected lottery reward and the status quo are far apart, a small perception change does not affect the choice and thus has no impact on outcomes. A marginal change in perception is consequential only when the two alternatives are perceived to be a *tie*: that is, when $qr_1 + (1 - q)r_2 = s$. The design of the optimal perception strategy, then, must condition on a tie occurring.

Conditioning on a tie tends to increase the weight placed on more extreme probabilities because perceptions close to 0 or 1 are more likely to lead to a tie than are perceptions close to 1/2. To see this, consider the following two lotteries, labelled with their perceived probabilities:



The perceived value of the first lottery is $(r_1 + r_2)/2$. Ex ante, before the rewards r_1 and r_2 are realized, the value of this lottery is normally distributed with mean 0 and variance 1/2. The perceived value of the second lottery is r_2 . Ex ante, the second lottery also has mean 0, but it has a higher variance (equal to 1). When the status quo is high, the higher variance makes a tie with the second lottery more likely than with the first.

More generally, for any given q , the perceived expected reward from the lottery, $qr_1 + (1 - q)r_2$, is normally distributed with mean 0 and variance $q^2 + (1 - q)^2$. Conditional on q , the likelihood that the agent perceives a tie is $\phi_q(s)$, where ϕ_q is the density for the normal distribution $N(0, q^2 + (1 - q)^2)$. Viewed as a function of q and suppressing s from the notation, we define the *weighting function* $w(q)$ to be equal to $\phi_q(s)$. Higher values of $w(q)$ correspond to increases in the expected fitness loss caused by an erroneous perception q of the objective probability p .⁷ When $s > 1$, the weight $w(q)$ is U-shaped, as depicted in Figure 3A.

⁷In the next subsection, where we derive the optimal perception strategy, we find that the correct weight given to various values of q differs from $w(q)$ by a factor that does not affect the direction of the distortions.




Figure 3: The likelihood of a tie between a lottery and the status quo as a function of the perceived probability for $s = 2$.

How should the agent distort probabilities in light of the U-shaped weighting function? We show that increasing the steepness of the perception function tends to reduce the effect of errors in perception. One can view this as focusing greater attention on probabilities at which the perception is steeper.⁸ For a U-shaped weighting function, more attention should be focused on extreme probabilities than on intermediate ones, suggesting that probabilities should be distorted according to an inverted S-shape, as in Figure 2. In the next subsection, we clarify the trade-off between attentiveness and correctness of perception that determines the optimal perception strategy.

Alternatively, the optimal distortion can be understood by an analogy to the winner's curse. Consider the naïve perception strategy $m(p) \equiv p$. For each p , that strategy leads to unbiased perception of the expected lottery reward in the sense that

$$E[\tilde{r} - r \mid p] = 0,$$

where $r = pr_1 + (1-p)r_2$ and $\tilde{r} = qr_1 + (1-q)r_2$ are, respectively, the true and the perceived expected rewards from the lottery. Although the perception is unbiased unconditionally, it *is* biased conditional on the agent perceiving a tie between the lottery and the status quo. In particular, when the status quo is high, one can show that

$$E[\tilde{r} - r \mid p \text{ and } \tilde{r} = s] > 0.$$

⁸See Tversky and Kahneman (1992) for a similar interpretation.

In case of a tie, the naïve perception strategy tends to overvalue the lottery because equality with s is more likely to occur if the error ε increases the perceived value of the lottery than if it decreases it.

Relative to the naïve strategy, the optimal strategy decreases the perceived value of the lottery conditional on a tie (this is loosely analogous to bid-shading in common value auctions). It turns out that exaggerating small probabilities (and underreporting large ones) does exactly that. To see how, consider the typical structure of lotteries that the agent perceives as a tie when the status quo is large. One possibility is that the higher probability branch is associated with a large reward, while the lower probability branch has a smaller reward, as in the flight lottery described in the Introduction. Alternatively, as in some casino games, the lottery can have a low probability of a very large reward coupled with a higher probability of a lower reward. Lotteries like the flight lottery are much more common because very large rewards are rare. In case of a tie, reducing the perception of high probabilities and exaggerating small ones therefore tends to reduce the perceived value of the lottery, helping to overcome the winner’s curse.

3.2 Outline of proof

To make the result as transparent as possible, in this subsection we outline a direct proof of Theorem 1 (as opposed to proving it as a corollary of the analogous result for the general model).

Fix s . We begin by identifying those decision problems in which, for a given perception strategy m satisfying $m(p) \in [p - \sigma, p + \sigma]$,⁹ the agent may end up choosing suboptimally. Intuitively, that will occur when the value of the lottery is close to the status quo.

For $r_1 \neq r_2$, let $p^*(r_1, r_2)$ be the solution to $pr_1 + (1 - p)r_2 = s$; that is, given r_1 and r_2 , the lottery is the optimal choice whenever p lies on one side of the threshold $p^*(r_1, r_2)$, and the status quo is optimal on the other side (which side depends on which of r_1 or r_2 is greater). A decision problem is difficult in the sense that the agent can choose suboptimally if the parameters (p, p^*) lie in the set

$$D = \{(p, p^*) : p^* \in [m(p) - \sigma, m(p) + \sigma]\}.$$

To see this, consider p^* outside of $[m(p) - \sigma, m(p) + \sigma]$. Since p is within that interval,

⁹Messages outside of $[p - \sigma, p + \sigma]$ are never optimal.




Figure 4: The set D of parameters at which a suboptimal choice may occur. An inverted S-shaped perception strategy m makes D narrower at values of p^* farther from $1/2$ (such as p_2^*) at the expense of making it wider at values close to $1/2$ (such as p_1^*).

$q \in \{m(p) - \sigma, m(p) + \sigma\}$, and hence q and p lie on the same side of the threshold p^* , implying that the choice based on q is optimal. Figure 4 illustrates the set D .

Given a strategy $m(\cdot)$, define the *ex ante* expected loss

$$L = E[\max\{pr_1 + (1-p)r_2, s\} - f(\ell, m(p) + \varepsilon)],$$

where the expectation is over the lottery ℓ and the noise ε . The loss L measures how much the agent's expected reward $f(\ell, q)$ falls below the first-best that can be attained in the absence of noise. The following lemma expresses the loss L as a weighted integral over the set D . Define the weighting function $\pi(p) = \lambda(p)w(p)$, where $w(p)$ is the likelihood of a tie as defined in the last subsection, and $\lambda(p) = E[(r_1 - r_2)^2 \mid pr_1 + (1-p)r_2 = s]$.

Lemma 1. *The expected loss satisfies $L = \frac{1}{2} \int_D |p^* - p| \psi(p) \pi(p^*) dp dp^*$.*

Given a threshold probability p^* , the agent suffers a large loss when (i) p^* is likely to generate a tie, and (ii) r_1 and r_2 tend to be far apart, making the value of the lottery sensitive to the probabilities. The first effect is captured by the $w(p^*)$ term, and the second by $\lambda(p^*)$. Combining the effects, the lemma indicates that the loss tends to be small when

the set D is both narrow and adheres closely to the diagonal.

Figure 4 illustrates how the slope of the perception strategy affects the loss L . The set D is narrow precisely when $m(p)$ is steep. Thus the inverted S-shaped perception strategy depicted in Figure 4 performs well toward the extremes at the expense of poorer performance at intermediate probabilities. If perception errors at intermediate probabilities generate smaller losses than those at more extreme probabilities then this leads to an overall gain. The following lemma confirms that this is indeed the case.

Lemma 2. *If $s > 3^{1/4}$, then the weighting function $\pi(p)$ is U-shaped: it is decreasing for $p < 1/2$, increasing for $p > 1/2$, and symmetric with respect to $p = 1/2$.*

To derive the optimal perception strategy, note that the integral in Lemma 1 can be minimized pointwise with respect to p . Thus, for each p , the optimal message satisfies

$$m^*(p) \in \arg \max_m \int_{m-\sigma}^{m+\sigma} |p^* - p| \pi(p^*) dp^*.$$

Taking the first order condition with respect to m gives the following characterization of the optimal strategy.

Lemma 3. *The optimal perception error $q - p$, weighted by $\pi(q)$, is unbiased; that is,*

$$E[(p - q)\pi(q)] = \sum_{\varepsilon \in \{-\sigma, \sigma\}} (m^*(p) + \varepsilon - p)\pi(m^*(p) + \varepsilon) = 0. \quad (2)$$

Theorem 1 follows from Lemmas 2 and 3. To see that the U-shaped weight implies that it is optimal to exaggerate small probabilities, consider $p < 1/2$ and $s > 3^{1/4}$. Suppose the observation center sends the unbiased message $m = p$, so that the perception is either $p - \sigma$ or $p + \sigma$. A marginal increase in m increases the loss by $\pi(p + \sigma)$ if the error is σ , and decreases the loss by $\pi(p - \sigma)$ if the error is $-\sigma$. Since $\pi(p - \sigma) > \pi(p + \sigma)$, increasing the message reduces the expected loss.

By symmetry, the optimal perception $m^*(p)$ at $-s$ is identical to that at s . The agent therefore exhibits an inverted S-shaped perception bias both when his status quo is high and when it is low relative to typical rewards in his environment. We focus on the first case because of Proposition 1.

4 The general case

In this section, we return to the general model from Section 2. Compared to the special case of Section 3, we now allow for a general distribution of rewards $\rho(r_1, r_2)$ with finite third moments, and for general perception formation $q = c(m, \varepsilon)$, where c is continuous, increasing in m , and continuously differentiable in m . The perception is nontrivially stochastic in the sense that for every m and q , $\Pr(c(m, \varepsilon) = q) < 1$.

The additional generality in the perception formation demonstrates that the pattern of distortions identified in the special case is not driven by the naïveté of the decision center. In Section 3, the decision center interprets the received message $m + \varepsilon$ at face value, failing to take into account the messaging strategy $m(\cdot)$ employed by the observation center. The general model allows for (but does not require) a decision center that, in equilibrium, correctly interprets the message, taking into account how the observation center codes the probability. See Section 4.1 for a detailed example.

In the general formulation of the model, messages are no longer directly comparable to probabilities. Instead, we compare the optimal perception under two objectives. In the *reward maximization problem*, defined in (1), the optimal strategy $m^*(p; s)$ maximizes the agent’s ex ante expected reward. We use as a benchmark the *precision maximization problem*, in which the optimal strategy $\hat{m}(p)$ minimizes the mean square error in perception;¹⁰ that is, for each p ,

$$\hat{m}(p) \in \arg \min_m E [(c(m, \varepsilon) - p)^2].$$

The precision-maximizing perception is a natural generalization of the unbiased perception strategy that we use as the benchmark in Section 3: when noise in communication is additive, the mean square error is minimized by unbiased perception.¹¹

The analysis in Section 3 makes use of the weight $\pi(p)$ that measures the importance of precise perception at each probability p . We construct a general weighting function as follows. Without loss of generality, we normalize $E[(r_1 - r_2)^2]$ to 1, and define a density $\tilde{\rho}(r_1, r_2) = (r_1 - r_2)^2 \rho(r_1, r_2)$. For each $p \in [0, 1]$, let $r(p) = pr_1 + (1 - p)r_2$, where the pair (r_1, r_2) is drawn according to $\tilde{\rho}$, and let $d_p(\cdot)$ be the density of $r(p)$. The weighting function is defined to be $\pi(p; s) = d_p(s)$.

The weight $\pi(p; s)$ captures how often the lottery with probability parameter p ties with

¹⁰Solutions to the two problems are guaranteed to exist, but there may be multiple optima, making the functions m^* and \hat{m} not uniquely defined. Our results hold for any optimal pair of m^* and \hat{m} .

¹¹In a related setting, Woodford (2012a,b) takes precision maximization as the objective of perception.

the status quo under the density $\tilde{\rho}$. Precise perception is relatively important at lotteries that have a large difference between the rewards, which is reflected in the additional mass assigned to those lotteries by $\tilde{\rho}$ relative to ρ . Notice that this definition of the weight coincides with the one from Section 3.

For each message m , let $\bar{q}(m) = \sup_{\varepsilon} c(m, \varepsilon)$ and $\underline{q}(m) = \inf_{\varepsilon} c(m, \varepsilon)$ denote the most extreme perceptions. We require the following technical assumptions:

A1 For each m , $\bar{q}(m) - \underline{q}(m) \leq \sigma$, where $\sigma \in (0, 1/4)$.

A2 The extremes cover the full range from 0 to 1, that is, $\underline{q}(\underline{m}) = 0$ and $\bar{q}(\bar{m}) = 1$.

A3 There exists a finite upper bound r^* such that for any $p_1 \neq p_2$, the densities of the random variables $r(p_1)$ and $r(p_2)$ do not intersect above r^* .

Assumption A3 is a regularity condition ensuring that the tails of the densities of $r(p)$ are well ordered across different values of p . The condition rules out densities that, for some pair p_1 and p_2 , alternate infinitely often to the right of any point. In addition, it requires that there is an upper bound on intersections that is uniform across p_1 and p_2 .

Define the reward-maximizing perception $q^* = c(m^*(p; s), \varepsilon)$ and the precision-maximizing perception $\hat{q} = c(\hat{m}(p), \varepsilon)$.

Theorem 2. *For any status quo s , the optimal message function $m^*(p; s)$ is nondecreasing in p . Furthermore, if $s > r^*$ and $p \in [0, 1] \setminus (1/2 - \sigma, 1/2 + \sigma)$, then $|m^*(p; s) - 1/2| \leq |\hat{m}(p) - 1/2|$, with the inequality being strict if $p \in (\sigma, 1 - \sigma)$.*

Proposition 1 suggests that, in the long-run, the status quo is likely to exceed r^* . In that case, the theorem indicates that, relative to the precision-maximizing perception, the optimal perception of small probabilities is biased upward and that of large probabilities is biased downward.

The remainder of this section outlines two lemmas that form the main steps in the proof of Theorem 2. The next lemma shows that the first order conditions of the two optimization problems differ only in the weight attributed to various perception errors. Let $c_m = \frac{\partial c}{\partial m}$.

Lemma 4. *For any $p \in [0, 1]$ such that $m^*(p), \hat{m}(p) \in (\underline{m}, \bar{m})$,*

$$E[\pi(q^*)(p - q^*)c_m(m^*, \varepsilon)] = 0,$$

and $E[(p - \hat{q})c_m(\hat{m}, \varepsilon)] = 0.$

As in Section 3, the weighting function is U-shaped.

Lemma 5. *For all $s > r^*$, $\pi(p; s)$ is decreasing for $p < 1/2$ and increasing for $p > 1/2$.*

The last lemma generalizes the intuition from Section 3.1 based on the observation that, when the status quo s is large, extreme probabilities are more likely to generate ties with s . In that section, we show that the density of the expected reward $r(p)$ becomes more spread out as p moves farther from $1/2$. That observation turns out to hold in general: we show in the appendix that when $|p_1 - 1/2| > |p_2 - 1/2|$, $r(p_1)$ is a mean-preserving spread of $r(p_2)$. Combined with the regularity condition A3, that implies that $r(p_1)$ has a thicker tail than $r(p_2)$, making $r(p_1)$ more likely to tie with a high status quo, which in turn implies the lemma and hence Theorem 2.

4.1 Equilibrium perception

The general model focuses on optimization at the observation stage, fixing the behavior of the decision center. One might expect that evolutionary pressures should also lead to optimization at the decision stage. We now present a simple example to illustrate how our model can accommodate optimization at both stages.¹² For the purpose of the example, we assume that the optimization at the decision stage is unconstrained; in the general model, by not requiring optimization at the decision stage, we are implicitly allowing for additional constraints at that stage.

The function $c(m, \varepsilon)$ that identifies the perception q captures not only the physical properties of the communication channel between the two centers, but also the way in which the decision center decodes the stimulus it receives. For optimization by the observation center, only the combination of those two elements matters. Here we separate them and explicitly model both the noise and the interpretation of the stimulus.

The objective probability p is uniformly drawn from $\{0, 1/4, 1/2, 3/4, 1\}$. The rewards r_1 and r_2 are drawn from the standard normal distribution, and all three lottery parameters are independent. For simplicity, the value of the status quo is fixed at 1.5. Upon observing p , the observation center sends a message $m \in [0, 1]$ that leads to a stimulus $v(m, \varepsilon)$ arriving at the decision center, with noise ε uniformly drawn from $[0, 1]$. The mapping $v : [0, 1]^2 \rightarrow \{0, 1/2, 1\}$, depicted in Figure 5, describes the communication technology between the two centers. Messages close to 0 are likely to result in stimulus 0, messages

¹²To keep the analysis simple, some details of the example deviate from the main model.




Figure 5: The communication technology: a message $m \in [0, 1/2]$ results in stimulus $v = 0$ with probability $1 - 2m$, and in stimulus $1/2$ with probability $2m$. Messages above $1/2$ similarly lead to stimuli $1/2$ or 1 .

close to $1/2$ are likely to result in stimulus $1/2$, and messages close to 1 are likely to result in stimulus 1 . Given any stimulus v , the decision center forms a perception $q(v)$ of the probability p , and chooses the lottery if and only if $r_1 q(v) + r_2(1 - q(v)) \geq s$. Restricting to a decision strategy that is linear in $q(v)$ is without loss of generality: the optimal choice as a function of all of the information available to the decision center is guaranteed to take this form. In terms of the notation in the main model, $c = q \circ v$.

An *equilibrium perception* consists of a message strategy $m(p)$ together with an interpretation strategy $q(v)$ such that (i) $q(v)$ maximizes the expected reward from the chosen action given $m(\cdot)$, and (ii) $m(p)$ maximizes the expected reward from the chosen action given $q(\cdot)$. Condition (i) is equivalent to $q(v) = E[p \mid v(m(p), \varepsilon) = v]$ given $m(\cdot)$.

To find an equilibrium, consider a message strategy of the form $m(0) = 0$, $m(1/4) = x$, $m(1/2) = 1/2$, $m(3/4) = 1 - x$, and $m(1) = 1$, with $x \in (0, 1/2)$. By symmetry, the equilibrium interpretation strategy satisfies $q(1/2) = 1/2$, and, by Bayes' rule, $q(0; x) = \frac{1-2x}{8(1-x)}$.

The equilibrium value of x is determined by an indifference condition. For $p = 1/4$, the observation center must be indifferent between the stimuli 0 and $1/2$, because otherwise it

can improve its payoff by changing x . Thus,

$$\int_{q(0;x)}^{\frac{1}{4}} \pi(p^*) \left(\frac{1}{4} - p^* \right) dp^* = \int_{\frac{1}{4}}^{\frac{1}{2}} \pi(p^*) \left(p^* - \frac{1}{4} \right) dp^*,$$

since the left-hand side is the loss due to incorrect perception when the decision center receives stimulus 0 and forms perception $q(0; x)$, and the right-hand side is the loss when the realized perception is $q(1/2) = 1/2$. The indifference condition is solved by $x \doteq 0.24$, implying that the perception $q(0; x)$ formed upon receiving stimulus 0 is approximately 0.09. By symmetry, $x \doteq 0.24$ solves the indifference condition at $p = 3/4$. It is easy to verify that the message strategy is optimal at $p = 0, 1/2$, and 1.

The equilibrium strategy maximizes the agent's fitness but does not maximize the precision of perception. To see that, consider $p = 1/4$. The agent's strategy results in perception $q = 0.09$ or $q = 1/2$. By sending the message 0, the observation center could guarantee a perception of $q = 0.09$. Since the perception error $1/4 - .09$ is smaller than the error $1/2 - 1/4$, this change would increase precision.

For the given strategy of the decision center, the observation center exaggerates the probability when $p = 1/4$ compared to the precision-maximization benchmark. As in our main model, the difference between fitness and precision maximization arises from the fact that positive and negative errors result in different likelihoods that the agent perceives the decision problem as a tie. This difference arises independently of whether the decision center interprets the messages optimally, which is why we focus on optimization at the observation stage.

5 Other biases

This section provides a sketch of how two other behavioral biases—an S-shaped value function and a status quo bias—can arise as optimal responses to noise in information processing. Although certain details of the examples in this section differ from the main model, the intuition remains the same: in each case, relative to unbiased perception, the behavioral bias mitigates errors arising from the failure to condition on ties.

The main change from the rest of the paper is that we now focus on distortions in the evaluation of rewards rather than probabilities. The agent chooses between a status quo s and an alternative action that pays a (certain) reward r . As before, there are two stages

of decision making. First, the observation center learns the value of r and sends a message m to the decision center. The decision center receives the message with noise, forms a perception $v = m + \varepsilon$ of r , and chooses the alternative action if and only if $v > s$. Given r , v , and s , the agent receives reward

$$f(r, v, s) = \begin{cases} r & \text{if } v > s, \\ s & \text{otherwise.} \end{cases}$$

To keep the problem nontrivial, we assume (as in the main model) that the observation center does not have all of the information needed to determine the optimal action. To capture this, we assume that, for some fixed \bar{s} , the status quo s is distributed as $N(\bar{s}, 1)$, and that the perception strategy depends only on \bar{s} , not on s itself.¹³ In addition, we assume that $r \sim N(0, 1)$ and ε attains values σ or $-\sigma$ with equal probabilities. All random variables are independent.

We consider two related behavioral biases that can arise depending on the space of possible perception strategies.

5.1 S-shaped value function

We first suppose that the observation center can choose an arbitrary strategy $m(r; \bar{s})$ depending on r and \bar{s} . We find that the optimal perception exhibits a key feature of loss aversion, namely, convexity of perceived rewards for losses (relative to \bar{s}) and concavity for gains.

The optimal perception strategy is defined by

$$m^*(r; \bar{s}) \in \arg \max_{m \in \mathbb{R}} E[f(r, m + \varepsilon, s)],$$

where the expectation is over ε and s .

It is straightforward to verify that the optimal perception satisfies the first-order condition

$$E[\phi(\bar{s} - v)(v - r)] \equiv E[\phi(\bar{s} - m^*(r; \bar{s}) - \varepsilon)(m^*(r; \bar{s}) + \varepsilon - r)] = 0, \quad (3)$$

where the expectation is with respect to the noise ε , and ϕ is the standard normal density.

¹³If the observation center could condition its strategy on the exact value of s , the first-best choice could be trivially obtained by allowing the observation center to compare the two options and send a strong signal indicating which is preferred.




Figure 6: Optimal perception of the reward r (solid curve) relative to unbiased perception $m(r) = r$ (dashed line).

The first-order condition requires the perception to be unbiased conditional on a tie; (3) is equivalent to $E[v - r \mid v = s] = 0$. The term $\phi(\bar{s} - v)$ is the likelihood that, given \bar{s} and v , the status quo s attains the value that leads to a tie. The optimal stochastic perception is unbiased with respect to this weight. Figure 6 depicts the solution of (3). The optimal perception is S-shaped. Relative to naïve, unbiased perception, the optimal perception is more attentive to rewards close to the average status quo, which are the ones that most often lead to a tie.

Proposition 2. *The optimal perception satisfies $m^*(r; \bar{s}) = r + \mu(r - \bar{s})$, where μ is positive and concave above 0, negative and convex below 0, and $\mu(-x) \equiv -\mu(x)$.*

The proposition follows directly from (3); we omit the details of the proof.

5.2 Status quo bias

We now reexamine the same perception problem except that we restrict the strategy space: the message function $m(r; \bar{s})$ must take the form $r - e(\bar{s})$ for some $e(\bar{s}) \in \mathbb{R}$. We interpret the term $e(\bar{s})$ as a measure of the status quo bias. Nature chooses the optimal bias across all draws of the reward r . Accordingly, the optimal strategy satisfies

$$e^*(\bar{s}) \in \arg \max_{e \in \mathbb{R}} E[f(r, r - e + \varepsilon, s)],$$

where the expectation is over r , ε , and s .

In this case, the optimal perception solves the first-order condition

$$E \left[\phi \left(\frac{\bar{s} + e^*(\bar{s}) - \varepsilon}{\sqrt{2}} \right) (\varepsilon - e^*(\bar{s})) \right] = 0. \quad (4)$$

Again, the condition is equivalent to requiring perception to be unbiased conditional on a tie. The first term in the expectation is the likelihood that $r - s$ is such that a tie arises between $v = r - e^*(\bar{s}) + \varepsilon$ and s .

The following proposition directly follows from (4).

Proposition 3. *If $\bar{s} > 0$, then the status quo bias $e^*(\bar{s})$ is positive; relative to unbiased perception, the agent overvalues the status quo.*

Suppose $\bar{s} > 0$. If the agent uses the naïve, unbiased perception strategy, then a positive value of ε leads to a tie more often than a negative value does. The unbiased strategy therefore exhibits a kind of winner’s curse that the positive status quo bias mitigates. By symmetry, the optimal status quo bias is the opposite when the average status quo is negative; that is, $e^*(\bar{s}) = -e^*(-\bar{s})$. Since an optimizing agent is likely to have a status quo that exceeds the average reward, a positive status quo bias should be typical, but a negative bias may arise in adverse circumstances.

6 Discussion of debiasing

The gap between the normative basis of expected utility theory and the descriptive origins of prospect theory (Thaler, 2000) has spurred an ongoing debate on “debiasing”. As Fischhoff (1982) writes:

Once judgmental biases are identified, researchers start trying to eliminate them using one of two strategies. The first accepts the existence of the bias and concentrates on devising schemes, such as training programs, that will reduce it.

Jolls and Sunstein (2006) argue that many laws are designed to counteract behavioral biases.

This paper offers a normative foundation for prospect theory biases as optimal responses to constraints in information processing. The optimizing role of biased perception suggests

that caution is warranted when considering whether deviations from expected utility should be eliminated; removing biases across all decision problems would harm the decision-maker.

Our theory suggests that probability biases are helpful in certain settings. For example, Barseghyan et al. (2013) document overvaluation of small probabilities using data on insurance deductible choices; clients who overweight low-probability losses prefer smaller deductibles than would unbiased decision-makers. De Giorgi and Legg (2012) explain the equity premium puzzle by pointing out that agents with prospect theory preferences overvalue the probability of rare market crashes. If errorless perception were possible, probability biases would be harmful in these cases (relative to correct perception). If, however, perception is noisy, then the observed biases can be beneficial in problems with a low probability of generating a loss. Overvaluation of small probabilities in those problems can be understood as a kind of cautiousness, without which the agent would select the risky action too often when perception errors reduce the perceived likelihood of a loss.

This is not to say that decision making cannot be improved upon. Using horse-race data, Thaler and Ziemba (1988) and Snowberg and Wolfers (2010) document excessive betting on low probability events that pay large rewards, leading to expected losses. While our theory suggests that overweighting of small probabilities can be helpful overall, it can also be harmful in settings where the distribution of lotteries differs from that faced by the decision-maker across all problems.

Why then does overvaluation of small probabilities in settings like the racetrack persist? When probabilities and rewards are processed separately, Nature must design the optimal bias across all types of problems. When high rewards are rare and the status quo relatively good, lotteries that tie with the status quo typically feature a high probability gain; otherwise, the low probability of a gain must be compensated by a very high reward. From the ex ante perspective, the distortion is optimal; ex post, it can be harmful. Put differently, debiasing may be beneficial in certain circumstances, but only in those that, from an evolutionary perspective, rarely result in a tie.

7 Open questions

In this paper, we examine noise in evaluation of probabilities and rewards in separate models. When only probabilities are distorted, we obtain inverted S-shaped probability weighting, and when only rewards are distorted, we obtain an S-shaped value function and the status quo bias. Combining noise in both parameters within a single model may lead

to interesting interaction effects but appears to be intractable. Since Herold and Netzer (2010) show that inverted S-shaped probability weights are an optimal response to an S-shaped value function, we conjecture that a combined model would strengthen the degree of probability weighting.

Our analysis precludes the possibility that perception of probabilities depends on the value of rewards. At the other extreme, where perception can depend on the exact realized rewards, the first-best can be attained. But what if perception depends on imperfect information about realized rewards? We conjecture that this would introduce a difference between perception of probabilities of gains and those of losses. In addition to the inverted S-shape, the agent would tend to put less weight on the probability of gains (and more on the probability of losses) since doing so would help to correct for the winner’s curse. This conjecture is broadly consistent with experimental results on reference-dependent probability weights (Tversky and Kahneman, 1992).

Our restriction to binary prospects is again due to issues of tractability. Allowing for more outcomes leads to interdependencies in the optimal perception that significantly complicate the problem. We expect that optimal perception in the general case would share some features of cumulative prospect theory, with inverted S-shaped probability weights that depend on the relative magnitudes of the probabilities in the lottery.

A Proof of Proposition 1

Lemma 6. *For every s_t , $\Pr(s_{t+1} > s_t \mid s_t) \geq \Pr(s_{t+1} < s_t \mid s_t)$.*

Proof of Lemma 6. If both rewards in the lottery are below s_t , then $s_{t+1} = s_t$. Similarly, if both rewards are above s_t , then $s_{t+1} > s_t$. Thus it suffices to prove that the claim holds conditional on r_{1t} and r_{2t} lying on opposite sides of s_t . Consider $p > 1/2$ and reward realizations $(r_{1t}, r_{2t}) = (r_1, r_2)$ such that $r_1 > s_t > r_2$. By the symmetry of the reward distribution and the distribution of p , it suffices to prove that

$$\begin{aligned} & \Pr(s_{t+1} > s_t \mid \ell_t = \ell) + \Pr(s_{t+1} > s_t \mid \ell_t = \ell') \\ & \geq \Pr(s_{t+1} < s_t \mid \ell_t = \ell) + \Pr(s_{t+1} < s_t \mid \ell_t = \ell'), \quad (5) \end{aligned}$$

where $\ell = (r_1, p; r_2, 1 - p)$ and $\ell' = (r_1, 1 - p; r_2, p)$. Letting

$$a = \Pr(q_t r_1 + (1 - q_t)r_2 > s_t \mid p_t = p)$$

and $b = \Pr(q_t r_1 + (1 - q_t)r_2 > s_t \mid p_t = 1 - p),$

we get

$$\Pr(s_{t+1} > s_t \mid \ell_t = \ell) + \Pr(s_{t+1} > s_t \mid \ell_t = \ell') = pa + (1 - p)b$$

and $\Pr(s_{t+1} < s_t \mid \ell_t = \ell) + \Pr(s_{t+1} < s_t \mid \ell_t = \ell') = (1 - p)a + pb.$

Monotonicity of the perception q_t implies that $a \geq b$. Since $p > 1/2$ by assumption, (5) follows. \square

Proof of Proposition 1. We will establish the following induction step. If for some $\xi \in (0, 1]$ and all $s^* \in \mathbb{R}$, $\limsup_t \Pr(s_t < s^*) \leq \xi$ then for all $s^{**} \in \mathbb{R}$, $\limsup_t \Pr(s_t < s^{**}) \leq \xi/2$. The induction hypothesis holds for $\xi = 1$, and thus the induction step suffices to prove the result.

Choose any s^{**} and s^* such that $s^* > s^{**}$. Let

$$\tau = \Pr(s_{t+1} > s^{**} \mid s_t = s^{**})$$

and $\tau' = \Pr(s_{t+1} < s^{**} \mid s_t = s^{**}),$

and notice that $\tau \geq \tau' > 0$, where the first inequality follows from Lemma 6. Let $\xi' = \limsup_t \Pr(s_t < s^{**})$. We claim that

$$\xi' \leq \xi' - \tau\xi' + \tau'(\xi - \xi') + 1 - \tilde{\Phi}(s^*), \tag{6}$$

where $\tilde{\Phi}$ is the cumulative distribution of the perceived expected reward $qr_1 + (1 - q)r_2$ from the lottery. Since (6) holds for every s^* and $1 - \tilde{\Phi}(s^*) \rightarrow 0$ as $s^* \rightarrow \infty$, it must be that $\xi' \leq \xi' - \tau\xi' + \tau'(\xi - \xi')$, and therefore $\xi' \leq \frac{\tau'}{\tau + \tau'}\xi \leq \xi/2$.

It remains to prove (6). Let ξ'_t be the probability that $s_t \leq s^{**}$, and ξ_t the probability

that $s_t \leq s^*$. We have

$$\begin{aligned} \xi'_{t+1} &= \xi'_t - \Pr(s_{t+1} > s^{**} \mid s_t < s^{**})\xi'_t \\ &\quad + \Pr(s_{t+1} < s^{**} \mid s^{**} < s_t < s^*) (\xi_t - \xi'_t) \\ &\quad + \Pr(s_{t+1} < s^{**} \mid s_t > s^*) (1 - \xi_t). \end{aligned} \quad (7)$$

For $s_{t+1} \neq s_t$, let $t(s_t, s_{t+1})$ denote the density of the Markov kernel conditional on s_t . Noting that $t(s_t, s_{t+1})$ is non-increasing in s_t , it follows that

$$\begin{aligned} \Pr(s_{t+1} > s^{**} \mid s_t < s^{**}) &\geq \tau \\ \text{and} \quad \Pr(s_{t+1} < s^{**} \mid s^{**} < s_t < s^*) &\leq \tau'. \end{aligned}$$

In addition,

$$\Pr(s_{t+1} < s^{**} \mid s_t > s^*) (1 - \xi_t) \leq 1 - \tilde{\Phi}(s^*)$$

because the status quo makes a transition only if the perceived expected reward from the lottery exceeds the current status quo.

Substituting the three inequalities from the last paragraph into (7) gives

$$\xi'_{t+1} \leq \xi'_t - \tau\xi'_t + \tau'(\xi_t - \xi'_t) + 1 - \tilde{\Phi}(s^*).$$

Taking the limit supremum on both sides and noting that $1 - \tau - \tau' \geq 0$, we have

$$\begin{aligned} \xi' &\leq \limsup_t \left(\xi'_t - \tau\xi'_t + \tau'(\xi_t - \xi'_t) + 1 - \tilde{\Phi}(s^*) \right) \\ &\leq \xi' - \tau\xi' + \tau'(\xi - \xi') + 1 - \tilde{\Phi}(s^*), \end{aligned}$$

and therefore (6) holds. □

B Proofs for Section 3

The definition of the weight $\pi(q; s) = \lambda(q; s)w(q; s)$ is equivalent to

$$\pi(p; s) = \int_{-\infty}^{\infty} \Delta^2 \phi(s + \Delta(1-p)) \phi(s - \Delta p) d\Delta.$$

Proof of Lemma 1. Write the expected loss L as

$$\frac{1}{2} \int_{\sigma}^{1-\sigma} \int_{\{(r_1, r_2): (p, p^*(r_1, r_2)) \in D\}} |p^*(r_1, r_2) - p| |r_1 - r_2| \phi(r_1) \phi(r_2) dr_1 dr_2 \psi(p) dp.$$

For each p , the inner integral is over the set of pairs (r_1, r_2) for which a suboptimal choice can occur. When a suboptimal choice occurs, the difference in expected reward between the lottery and the status quo is $|p^*(r_1, r_2) - p| |r_1 - r_2|$. The factor $1/2$ reflects that the suboptimal choice occurs only for one of the two possible realizations of ε .

Consider the substitution $(p^*, \Delta) = (p^*(r_1, r_2), r_1 - r_2) = \left(\frac{s - r_2}{r_1 - r_2}, r_1 - r_2 \right)$ in the inner integral. It is straightforward to verify that the Jacobian associated with this substitution has determinant $\frac{1}{r_1 - r_2} = \frac{1}{\Delta}$. Therefore, the expression for L becomes

$$\frac{1}{2} \int_{\sigma}^{1-\sigma} \int_{m(p)-\sigma}^{m(p)+\sigma} |p^* - p| \pi(p^*) dp^* \psi(p) dp,$$

as needed. □

Proof of Lemma 2. First we compute $\lambda(q)$. Write the random variable $r_1 - r_2$ as $a\zeta + b\zeta^\perp$ where $\zeta = qr_1 + (1 - q)r_2$ and $\zeta^\perp = -(1 - q)r_1 + qr_2$. Note that ζ and ζ^\perp are independent. Comparing the coefficients, we obtain

$$a = -\frac{1 - 2q}{(1 - q)^2 + q^2}, \quad b = -\frac{1}{(1 - q)^2 + q^2}.$$

Conditional on $\zeta = s$, the random variable $r_1 - r_2$ has mean as and variance $V = b^2 \text{Var}(\zeta^\perp) = b^2((1 - q)^2 + q^2)$. Thus we have

$$\lambda(q) = E[(r_1 - r_2)^2 | \zeta = s] = (as)^2 + V = \frac{q^2 + (1 - q)^2 + s^2(1 - 2q)^2}{(q^2 + (1 - q)^2)^2}.$$

Multiplying the last expression by $w(q) = \phi_q(s)$ gives

$$\pi(q; s) = \frac{q^2 + (1 - q)^2 + s^2(1 - 2q)^2}{(q^2 + (1 - q)^2)^2} \frac{1}{\sqrt{q^2 + (1 - q)^2}} \phi_1\left(\frac{s}{\sqrt{q^2 + (1 - q)^2}}\right),$$

which is symmetric around $q = 1/2$.

Let $y = \frac{1}{\sqrt{q^2 + (1 - q)^2}}$ and note that y is increasing in q and attains values in $(1, \sqrt{2}]$ if

$q \in (0, 1/2]$. Therefore, it suffices to prove that

$$\pi(q(y); s) = \frac{1}{\sqrt{2\pi}} e^{-\frac{s^2 y^2}{2}} y^3 (1 + 2s^2 - s^2 y^2)$$

is decreasing in y on $(1, \sqrt{2}]$. Differentiating gives

$$\frac{\partial \pi(q(y); s)}{\partial y} = \frac{1}{\sqrt{2\pi}} e^{-\frac{s^2 y^2}{2}} y^2 [s^4 y^4 - 2(s^2 + 3)s^2 y^2 + 3(2s^2 + 1)].$$

This derivative is negative if the expression in the square brackets is negative, which is the case whenever

$$y^2 \in \left(\frac{s^2 + 3 - \sqrt{s^4 + 6}}{s^2}, \frac{s^2 + 3 + \sqrt{s^4 + 6}}{s^2} \right). \quad (8)$$

For $s > 3^{1/4}$, this interval contains $[1, 2]$, and therefore (8) holds for $y \in (1, \sqrt{2}]$. \square

Proof of Theorem 1. For $m \in [p - \sigma, p + \sigma]$, define the loss $l(m, p) = \int_{m-\sigma}^{m+\sigma} |p^* - p| \pi(p^*) dp^*$. Then $m^*(p)$ maximizes $-l(m, p)$. Notice that

$$\frac{\partial^2}{\partial p \partial m} (-l(m, p)) = \pi(m - \sigma) + \pi(m + \sigma) > 0,$$

and thus $m^*(p)$ is non-decreasing.

Consider $p \in [\sigma, 1/2]$. Suppose for contradiction that $m^* \leq p$ (here we suppress the argument of $m^*(p)$ from the notation). Then, by Lemma 2, $\pi(m^* - \sigma) > \pi(m^* + \sigma)$, and hence

$$\begin{aligned} E[\pi(m^* + \varepsilon)(m^* + \varepsilon - p)] &= \frac{1}{2}\pi(m^* + \sigma)(m^* + \sigma - p) + \frac{1}{2}\pi(m^* - \sigma)(m^* - \sigma - p) \\ &< \pi(m^* + \sigma) \left(\frac{1}{2}(m^* + \sigma - p) + \frac{1}{2}(m^* - \sigma - p) \right) \\ &= \pi(m^* + \sigma)(m^* - p) \leq 0, \end{aligned}$$

violating (2) in Lemma 3. Therefore, $m^* > p$. The argument for $p \in (1/2, 1 - \sigma]$ is analogous. \square

C Proofs for Section 4

Proof of Lemma 4. The second equation follows directly from the first-order condition for the precision maximization problem.

For the first equation, first note that

$$\pi(q; s) = \int_{-\infty}^{\infty} (r_1 - r_2(r_1; q))^2 \frac{\rho(r_1, r_2(r_1; q))}{(1 - q)} dr_1,$$

where $r_2(r_1; q)$ is the value of r_2 that leads to a tie given r_1 and q , that is, $r_2(r_1; q) = \frac{s - qr_1}{1 - q}$.

Let $F = E[f(\ell, c(m, \varepsilon)) \mid p]$ be the expected reward when the observation center observes p and sends message m . We have

$$F = E \left[\int_{c(m, \varepsilon)r_1 + (1 - c(m, \varepsilon))r_2 < s} s \rho(r_1, r_2) dr_1 dr_2 + \int_{c(m, \varepsilon)r_1 + (1 - c(m, \varepsilon))r_2 > s} (pr_1 + (1 - p)r_2) \rho(r_1, r_2) dr_1 dr_2 \right],$$

where the expectation is with respect to the noise ε , as are all other expectations for the remainder of this proof.

The expected reward F is continuous in m and the message space is compact. Therefore, an optimal message exists. Moreover, since F is continuously differentiable in m , the optimal message must, for each p , satisfy the first-order condition $\frac{\partial}{\partial m} F = 0$ whenever it is interior. Computing the derivative, we obtain

$$\begin{aligned} \frac{\partial}{\partial m} F &= E \left[c_m(m, \varepsilon) \int_{-\infty}^{\infty} (s - (pr_1 + (1 - p)r_2(r_1; c(m, \varepsilon)))) \frac{s - r_1}{(1 - c(m, \varepsilon))^2} \rho(r_1, r_2(r_1; c(m, \varepsilon))) dr_1 \right] \\ &= E \left[c_m(m, \varepsilon) \int_{-\infty}^{\infty} (c(m, \varepsilon)r_1 + (1 - c(m, \varepsilon))r_2(r_1; c(m, \varepsilon)) - (pr_1 + (1 - p)r_2(r_1; c(m, \varepsilon)))) \right. \\ &\quad \left. \cdot \frac{c(m, \varepsilon)r_1 + (1 - c(m, \varepsilon))r_2(r_1; c(m, \varepsilon)) - r_1}{(1 - c(m, \varepsilon))^2} \rho(r_1, r_2(r_1; c(m, \varepsilon))) dr_1 \right] \\ &= E \left[c_m(m, \varepsilon) \int_{-\infty}^{\infty} \frac{p - c(m, \varepsilon)}{1 - c(m, \varepsilon)} (r_1 - r_2(r_1; c(m, \varepsilon)))^2 \rho(r_1, r_2(r_1; c(m, \varepsilon))) dr_1 \right] \\ &= E [(p - c(m, \varepsilon))c_m(m, \varepsilon)\pi(c(m, \varepsilon))], \end{aligned}$$

as needed. □

The next two lemmas are required for the proof of Lemma 5.

Definition 3. Let X_1 and X_2 be real-valued random variables with distribution functions F_1 and F_2 , respectively. We say that X_1 is a *mean-preserving spread* of X_2 if

1. $E[X_1] = E[X_2]$, and
2. for every $x \in \mathbb{R}$,

$$\int_{-\infty}^x F_1(X)dX \geq \int_{-\infty}^x F_2(X)dX.^{14} \quad (9)$$

Lemma 7. Let X be a random variable with distribution function $F(\cdot)$. Let $Y_1(X)$ and $Y_2(X)$ be random variables with distribution functions $G_1(X)(\cdot)$ and $G_2(X)(\cdot)$, respectively, such that, for each X , $Y_2(X)$ is a mean-preserving spread of $Y_1(X)$. Finally, let Z_i be the random variable obtained by composition of X and Y_i (so that Z_i has distribution function $H_i(Z) = \int_{-\infty}^{\infty} G_i(X)(Z)dF(X)$). Then Z_2 is a mean-preserving spread of Z_1 .

Proof. We need to show that (i) Z_1 and Z_2 have the same mean, and (ii) $\int_{-\infty}^t H_2(Z)dZ \geq \int_{-\infty}^t H_1(Z)dZ$ for every t .

For claim (i),

$$E[Z_1] = \int_{-\infty}^{\infty} E[Y_1(X)]dF(X) = \int_{-\infty}^{\infty} E[Y_2(X)]dF(X) = E[Z_2].$$

For claim (ii),

$$\begin{aligned} \int_{-\infty}^t H_2(Z)dZ &= \int_{-\infty}^t \int_{-\infty}^{\infty} G_2(X)(Z)dF(X)dZ \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^t G_2(X)(Z)dZdF(X) \\ &\geq \int_{-\infty}^{\infty} \int_{-\infty}^t G_1(X)(Z)dZdF(X) \\ &= \int_{-\infty}^t \int_{-\infty}^{\infty} G_1(X)(Z)dF(X)dZ \\ &= \int_{-\infty}^t H_1(Z)dZ, \end{aligned}$$

¹⁴There is some disagreement in the literature on this terminology. Rothschild and Stiglitz (1970) and Müller (1998) use the term “mean-preserving spread” much more narrowly; what we call a mean-preserving spread, Müller (1998) calls a mean-preserving ϕ -increase in risk.

where the inequality follows from $Y_2(X)$ being a mean-preserving spread of $Y_1(X)$ for each X . \square

Recall that the random variable $r(p)$ is defined as $pr_1 + (1-p)r_2$, where the pair (r_1, r_2) is drawn according to the density $\tilde{\rho}(r_1, r_2) = (r_1 - r_2)^2 \rho(r_1, r_2)$.

Lemma 8. *If $|p_1 - 1/2| > |p_2 - 1/2|$ then $r(p_1)$ is a mean-preserving spread of $r(p_2)$.*

Proof of Lemma 8. For each r_1 and r_2 , define $r^+(p) = pr_1 + (1-p)r_2$, and $r^-(p) = (1-p)r_1 + pr_2$. Let $\hat{\ell}(r_1, r_2, p)$ be the binary lottery $(r^+(p), 1/2; r^-(p), 1/2)$. The symmetry of $\tilde{\rho}$ implies that $r(p)$ is equivalent in distribution to the compound lottery in which (r_1, r_2) is drawn from $\tilde{\rho}$, and then $r(p)$ is drawn from $\hat{\ell}(r_1, r_2, p)$.

By Lemma 7, it suffices to show that if $|p_1 - 1/2| > |p_2 - 1/2|$, then for each draw of (r_1, r_2) , the binary lottery $\hat{\ell}(r_1, r_2, p_1)$ is a mean-preserving spread of $\hat{\ell}(r_1, r_2, p_2)$. This is indeed the case since both lotteries have the same mean $(r_1 + r_2)/2$, and $r^+(p_2)$ and $r^-(p_2)$ lie in between $r^+(p_1)$ and $r^-(p_1)$. \square

Proof of Lemma 5. Consider p_1 and p_2 such that $|p_1 - 1/2| > |p_2 - 1/2|$. Let $X_1 = -r(p_1)$ and $X_2 = -r(p_2)$. By Lemma 8, X_1 is a mean-preserving spread of X_2 . Let $F_1(x) = \int_{-x}^{+\infty} d_{p_1}(x')dx'$ and $F_2(x) = \int_{-x}^{+\infty} d_{p_2}(x')dx'$ be the distribution functions of X_1 and X_2 , respectively.

By the regularity condition A3, either $d_{p_1}(x) > d_{p_2}(x)$ for all $x > r^*$, or $d_{p_1}(x) < d_{p_2}(x)$ for all $x > r^*$. For the sake of contradiction, suppose the latter. Then $F_1(x) < F_2(x)$ for all $x < -r^*$ and hence

$$\int_{-\infty}^{-r^*} F_1(x')dx' < \int_{-\infty}^{-r^*} F_2(x')dx',$$

which contradicts that X_1 is a mean preserving spread of X_2 . \square

Proof of Theorem 2. We first prove the second sentence of the proposition (the comparison between $m^*(p; s)$ and $\hat{m}(p; s)$). Consider $p \in [0, 1/2 - \sigma]$. The argument for $p > 1/2 + \sigma$ is analogous.

By Lemma 4 together with the corresponding inequalities for corner solutions,

$$E[(p - c(m^*, \varepsilon))c_m(m^*, \varepsilon)\pi(c(m^*, \varepsilon))] \begin{cases} \leq 0 & \text{if } m^* = \underline{m}, \\ = 0 & \text{if } m^* \in (\underline{m}, \overline{m}), \\ \geq 0 & \text{if } m^* = \overline{m}, \end{cases}$$

$$E[(p - c(\hat{m}, \varepsilon))c_m(\hat{m}, \varepsilon)] \begin{cases} \leq 0 & \text{if } \hat{m} = \underline{m}, \\ = 0 & \text{if } \hat{m} \in (\underline{m}, \overline{m}), \\ \geq 0 & \text{if } \hat{m} = \overline{m}. \end{cases}$$

Let $m_{1/2} = \sup\{m : \bar{q}(m) \leq 1/2\}$. For any $m > m_{1/2}$, it follows from assumption A1 that, for every ε , $c(m, \varepsilon) \geq \bar{q}(m) - \sigma \geq 1/2 - \sigma > p$. Hence m cannot satisfy the first-order condition for either problem, and we must have $\hat{m}(p), m^*(p) \leq m_{1/2}$. Therefore, we can restrict attention to messages m satisfying $m \leq m_{1/2}$. For this range, $\pi(c(m, \varepsilon))$ is decreasing in m .

Fixing p , let

$$\Delta^*(m', m) = E[f(\ell, c(m', \varepsilon)) - f(\ell, c(m, \varepsilon)) \mid p]$$

and

$$\hat{\Delta}(m', m) = \frac{\pi(p)}{2} \left(-\hat{L}(m', p) + \hat{L}(m, p) \right),$$

where $\hat{L}(m, p) = E[(c(m, \varepsilon) - p)^2]$. These two expressions may be interpreted as the payoff difference between messages m' and m under fitness and precision maximization, respectively, with the caveat that we have rescaled the payoffs in the precision maximization problem by $\frac{\pi(p)}{2}$.

We claim that

$$\Delta^*(m', m) > \hat{\Delta}(m', m)$$

whenever $m' \in (m, m_{1/2})$. This implies that $m^*(p) \geq \hat{m}(p)$ by Theorem 1 in Van Zandt (2002).

To prove the claim, first define

$$\delta^*(q', q) = E[f(\ell, q') - f(\ell, q) \mid p]$$

and

$$\hat{\delta}(q', q) = \frac{\pi(p)}{2} \left(-(q' - p)^2 + (q - p)^2 \right).$$

Note that $\Delta^*(m', m) = E[\delta^*(c(m', \varepsilon), c(m, \varepsilon))]$ and $\hat{\Delta}(m', m) = E[\hat{\delta}(c(m', \varepsilon), c(m, \varepsilon))]$. Since $c(m, \varepsilon)$ is increasing in m , it suffices to show that

$$\delta^*(q', q) > \hat{\delta}(q', q) \quad (10)$$

for all $q', q \in [0, 1/2]$ such that $q' > q$.

It is straightforward to verify that

$$\delta^*(q', q) = \int_q^{q'} (p - p^*)\pi(p^*)dp^*, \quad (11)$$

and

$$\hat{\delta}(q', q) = \int_q^{q'} (p - p^*)\pi(p)dp^*; \quad (12)$$

(11) follows from considering cases based on how q and q' compare to the critical probability p^* , while (12) can be verified directly by integration. Inequality (10) holds whenever $q < q' \leq 1/2$ because π is decreasing, and hence the integrand in (11) strictly exceeds that in (12) for all $p^* \in [q, q'] \setminus \{p\}$.

So far we have established only a weak inequality between $m^*(p)$ and $\hat{m}(p)$; we will show that the inequality must be strict for all $p \in (\sigma, 1/2 - \sigma)$. Note first that if $p > \sigma$ then, by assumptions A1 and A2, $c(\underline{m}, \varepsilon) \leq \sigma < p$ for all ε , and hence \underline{m} cannot solve the optimality condition for either problem. Similarly, \bar{m} cannot be optimal, and we can restrict attention to interior solutions.

Consider $p \in (\sigma, 1/2 - \sigma)$. Suppose for contradiction that $m^*(p) = \hat{m}(p) = m$, where $m \in (\underline{m}, \bar{m})$. By the first-order condition for the precision maximization problem,

$$E[(p - c(m, \varepsilon))c_m(m, \varepsilon)\pi(p)] = 0. \quad (13)$$

Since

$$(p - c(m, \varepsilon))c_m(m, \varepsilon)\pi(c(m, \varepsilon)) \geq (p - c(m, \varepsilon))c_m(m, \varepsilon)\pi(p) \quad (14)$$

for every ε , the left-hand side of (13) is less than $E[(p - c(m, \varepsilon))c_m(m, \varepsilon)\pi(c(m, \varepsilon))]$. Moreover, the inequality in (14) is strict unless $c(m, \varepsilon) = p$, which happens with probability less than 1. Therefore, $E[(p - c(m, \varepsilon))c_m(m, \varepsilon)\pi(c(m, \varepsilon))] > 0$, contradicting that m solves the reward maximization problem.

To prove that $m^*(p)$ is nondecreasing, note first that $\delta^*(q', q)$ is increasing in p when $q' >$

q . Since $\Delta^*(m', m; p) = E[\delta^*(c(m', \varepsilon), c(m, \varepsilon))]$ and $c(m, \varepsilon)$ is increasing in m , $\Delta^*(m', m)$ must be increasing in p whenever $m' > m$. The result then follows from Theorem 1 in Van Zandt (2002). \square

References

- Barseghyan, L., F. Molinari, T. O'Donoghue, and J. C. Teitelbaum (2013). The nature of risk preferences: Evidence from insurance choices. *American Economic Review* 103(6), 2499–2529.
- Berns, G. S. and E. Bell (2012). Striatal topography of probability and magnitude information for decisions under uncertainty. *Neuroimage* 59(4), 3166–3172.
- Bossaerts, P., K. Preuschoff, and M. Hsu (2009). The neurobiological foundations of valuation in human decision making under uncertainty. In P. W. Glimcher, C. F. Camerer, E. Fehr, and R. A. Poldrack (Eds.), *Neuroeconomics: Decision Making and the Brain*, pp. 353–365. Academic Press.
- Camerer, C., G. Loewenstein, and D. Prelec (2005). Neuroeconomics: How neuroscience can inform economics. *Journal of Economic Literature*, 9–64.
- Compte, O. and A. Postlewaite (2012). Cautiousness. Working Paper.
- De Giorgi, E. G. and S. Legg (2012). Dynamic portfolio choice and asset pricing with narrow framing and probability weighting. *Journal of Economic Dynamics and Control* 36(7), 951–972.
- Fischhoff, B. (1982). Debiasing. In D. Kahneman, P. Slavic, and A. Tversky (Eds.), *Judgment Under Uncertainty: Heuristics and Biases*, pp. 422–444. New York: Cambridge University Press.
- Frenkel, S., Y. Heller, and R. Teper (2012). Endowment as a blessing. Working Paper.
- Friedman, D. (1989). The S-shaped value function as a constrained optimum. *American Economic Review* 79(5), 1243–1248.
- Glimcher, P. W. (2005). Indeterminacy in brain and behavior. *Annual Review of Psychology* 56, 25–56.

- Glimcher, P. W. (2009). Choice: towards a standard back-pocket model. In P. W. Glimcher, C. F. Camerer, E. Fehr, and R. A. Poldrack (Eds.), *Neuroeconomics: Decision Making and the Brain*, pp. 503–521. Academic Press.
- Gul, F. and W. Pesendorfer (2008). The case for mindless economics. In A. Caplin and A. Schotter (Eds.), *The Foundations of Positive and Normative Economics*, pp. 3–39. Oxford University Press.
- Herold, F. and N. Netzer (2010). Probability weighting as evolutionary second-best. Working Paper.
- Jolls, C. and C. R. Sunstein (2006). Debiasing through law. *Journal of Legal Studies* 35, 199–241.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291.
- Kirkpatrick, L. A. and S. Epstein (1992). Cognitive-experiential self-theory and subjective probability: further evidence for two conceptual systems. *Journal of Personality and Social Psychology* 63(4), 534.
- McFadden, D. (1999). Rationality for economists? *Journal of Risk and Uncertainty* 19(1), 73–105.
- Müller, A. (1998). Comparing risks with unbounded distributions. *Journal of Mathematical Economics* 30(2), 229–239.
- Netzer, N. (2009). Evolution of time preferences and attitudes toward risk. *American Economic Review* 99(3), 937–955.
- Rayo, L. and G. S. Becker (2007). Evolutionary efficiency and happiness. *Journal of Political Economy* 115(2), 302–337.
- Robson, A. and L. Samuelson (2011). The evolution of decision and experienced utilities. *Theoretical Economics* 6(3), 311–339.
- Robson, A. J. (2001a). The biological basis of economic behavior. *Journal of Economic Literature* 39(1), 11–33.

- Robson, A. J. (2001b). Why would nature give individuals utility functions? *Journal of Political Economy* 109(4), 900–914.
- Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. a definition. *Journal of Economic Theory* 2(3), 225–243.
- Samuelson, L. and J. M. Swinkels (2006). Information, evolution and utility. *Theoretical Economics* 1(1), 119–142.
- Snowberg, E. and J. Wolfers (2010). Explaining the favorite–long shot bias: Is it risk-love or misperceptions? *Journal of Political Economy* 118(4), 723–746.
- Thaler, R. H. (2000). From homo economicus to homo sapiens. *Journal of Economic Perspectives* 14(1), 133–141.
- Thaler, R. H. and W. T. Ziemba (1988). Parimutuel betting markets: racetracks and lotteries. *Journal of Economic Perspectives* 2(2), 161–174.
- Tobler, P. N., G. I. Christopoulos, J. P. O’Doherty, R. J. Dolan, and W. Schultz (2008). Neuronal distortions of reward probability without choice. *Journal of Neuroscience* 28(45), 11703–11711.
- Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4), 297–323.
- Van Zandt, T. (2002). An introduction to monotone comparative statics. Available at <http://faculty.insead.edu/vanzandt/teaching/CompStatics.pdf>.
- Woodford, M. (2012a). Inattentive valuation and reference-dependent choice. Working Paper.
- Woodford, M. (2012b). Prospect theory as efficient perceptual distortion. *American Economic Review* 102(3), 41–46.