

Série des Documents de Travail

n° 2017-50
**Optimal Illusion of Control and Related
Perception Biases**

O. GOSSNER¹
J. STEINER²

Les documents de travail ne reflètent pas la position du CREST et n'engagent que leurs auteurs.
Working papers do not reflect the position of CREST but only the views of the authors.

¹ CREST; CNRS; Ecole polytechnique; Université Paris-Saclay; London School of Economics. E-mail : olivier.gossner@polytechnique.edu

² University of Edinburgh; Cerge-Ei. E-mail : Jakub.steiner@cerge-ei.cz

Optimal Illusion of Control and Related Perception Biases*

Olivier Gossner[†]

CREST, CNRS, École polytechnique, Université Paris-Saclay
and London School of Economics

Jakub Steiner[‡]

University of Edinburgh and Cerge-Ei

October 16, 2017

Abstract

We study perception biases arising under second-best perception strategies. An agent correctly observes a parameter that is payoff-relevant in many decision problems that she encounters in her environment but is unable to retain all the information until her decision. A designer of the decision process chooses a perception strategy that determines the distribution of the perception errors. If some information loss is unavoidable due to cognition constraints, then (under additional conditions) the optimal perception strategy exhibits the illusion of control, overconfidence, and optimism.

1 Introduction

Within economic discourse, the idea that the human mind is imperfect and that some information is necessarily lost during any decision process dates back at least to Simon

*We thank Andrew Clausen, Bard Harstad, Yuval Heller, David Hirshleifer, Ed Hopkins, Philippe Jéhiel, Drazen Prelec, József Sákovics, Colin Stewart, Andy Zapechelnnyuk, and the audiences at YEM 2016, BRIC 2016, GAMES 2016 and the Limited Cognition workshop in Barcelona for comments. Ludmila Matysková has provided excellent research assistance and Judith Levi has carefully edited the text. Steiner was supported by the Purkyne fellowship of the Czech Academy of Sciences and by the GACR grant 16-00703S.

[†]email: olivier.gossner@polytechnique.edu

[‡]email: jakub.steiner@ed.ac.uk

(1955). Under this premise, an agent faces, apart from the standard action choice, a problem of error management. In economic decision making, some mistakes are costlier than others, and this asymmetry has an impact on the frequency of different types of errors. In this paper, we relate error management to known observed behavioural biases.

For instance, the agent may not be able to perfectly retain all relevant information about her talent and thus may make decisions as under- or overconfident about her own talent. Or she may be partially informed about a state of nature without remembering perfectly how accurate her information is, thus under- or overestimating the precision of her information. Or she may face situations in which her payoffs depend more or less strongly on her own decisions, and imperfect assessment of these situations may lead to an illusion of control, or of lack of it.¹

This paper studies how and when behavioral biases such as overconfidence, overprecision, and illusion of control arise as the consequences of error management. In our model, the decision maker first receives some statistical information on a state of nature which can be either high or low, and this information translates into an objective belief p that the state is high. In a perception stage, she recalls this information imperfectly and ends up with a belief q that the state is high. Then, she faces a decision problem in which she chooses an action a from a fixed choice set and receives a utility $u(a, \theta)$ that depends both on the state θ and the action choice a . In this decision problem, the agent knows the utility function $u(\cdot, \cdot)$, and acts *as if* her belief about the state was q and hence her behavior may be distorted. At the perception stage, the agent does not know the utility function $u(\cdot, \cdot)$ but knows the distribution of utility functions that she will face at the decision stage. We study how the stochastic properties of the random decision problem translate into error management in the perception stage and into corresponding perception biases.

To this aim, we provide a definition of the loss associated with each misperception of q instead of p that builds on the standard economic welfare approach. The misperceived probability of the high state distorts choice and leads to a welfare loss in some of the encountered decision problems. We define the loss associated with this misperception as the welfare loss averaged across all decision problems in the agent's environment. We characterize the loss function, given the distribution of utility functions, and we report conditions under which the loss function exhibits useful comparative statics. In particular, we identify

¹*Illusion of control* refers to the overestimation of one's ability to impact one's own well-being. As an extreme example, casino visitors may overestimate the relevance of their choices over equivalent lotteries. The term was introduced by Langer (1975).

cases in which overconfidence generates a smaller loss than does under-confidence of a same magnitude.

Our main applied contribution relates such asymmetries of the loss function to the degree of control the agent has over her utility. We consider settings in which the agent has more control over her utility in the high state than in the low one, in the sense the distribution of utility functions is such that payoffs resulting from any action have a higher variance in the high state than in the low one. Then, for most of the realizations of the utility function observed by the agent, the payoff differences across actions are larger in the high state than in the low one. The agent's choice has a larger impact on her well-being in the high state than in the low one. We show that, under this assumption, an overvaluation of the objective probability p leads to a smaller loss than an undervaluation of the same magnitude.

Consider for instance the decision maker facing objective probability $1/2$. Intuitively, when payoff variance is higher in the high state than in the low state, her correct decision tends to be more correlated with the correct decision in the high state than in the low state. Over-focusing on the high state is therefore less detrimental in decision making than under-focusing on this state. Slightly more technically, a misperception of the objective probability p for a subjective probability q leads to a loss only in those decision problems in which the optimal actions at p and q differ. Thus, the loss arises if and only if the agent is indifferent between the two actions at some intermediate probability s in-between p and q . When an agent has more control in the high state than in the low one, the random expected payoff for any given action at probability s has a more spread-out density for higher values of the probability s , and ties between the two actions are less likely to arise in an upper neighborhood of the objective probability p than in the lower neighborhood of p of the same magnitude.

Therefore, overvaluing one's control over one's own well-being leads to a smaller loss than undervaluing one's control. When errorless perception is infeasible, optimization over the distributions of the perception errors leads to a second-best stochastic perception that is biased toward the state of the high control. The imperfect agent exhibits an illusion of control. We also show how this effect generates further behavioral phenomena. We provide applications in which the choice of highly talented or well-informed people has a larger impact on their well-being than the choice of less talented or less well-informed people, and show that an overconfidence bias is optimal in these cases when perfect perception is infeasible. The illusion of control can also generate optimism. An optimizing agent benefits

from a high spread of the payoffs to the available actions, and hence the bias toward the high-control state coincides with the bias toward the state in which the optimizing agent fares better.

Although cognition constraints are uncontroversial, little is known about their exact form. We examine three variants of the perception problem that differ in the set of the feasible stochastic perception strategies. In the first and simplest version, the agent observes the true probability p of the high state, chooses to memorize a probability m , and recalls probability $q = m + \varepsilon$, where ε is a small additive error of zero mean. We show that when the agent has more control in the high state than in the low one, the optimal perception strategy is biased upwards. In this case, the agent memorizes a probability m that exceeds the true probability p , and the recalled probability q exceeds p on average across the error ε . The agent in this setting is naive in the sense that she does not form Bayesian posteriors about the objective probability p upon recalling q ; instead, she treats the subjective probability q as if it were accurate. Such a model can capture an agent who makes mistakes in evaluating the available evidence and arrives at a perturbed posterior q instead of forming the correct posterior p . The optimal perception strategy optimizes the distribution of the updating errors.

In the second version of the perception problem, we consider a sophisticated agent who loses information during the decision process but rationally utilizes all the information she retains. The agent has a binary memory and thus forms a binary partition of the probability interval, remembers the cell of the partition into which the objective probability p belongs, and forms a Bayesian posterior about p at the decision stage. We show that, if the prior density of p is symmetric around $1/2$ and the agent has more control in the high state than in the low one, then the upper cell of the optimal partition is larger than the lower cell. Thus, the agent is more likely to hold a belief that attaches a high probability to the high state than a belief favoring the low state.

In our analysis of perception biases, we remain as agnostic as possible on the analytical formulation of perception problem and derive testable predictions on perception biases based on observed choice data. We introduce a model of stochastic perceptions that generalizes the previous two settings and assume that the agent chooses from a set of available stochastic perceptions, where each such perception is a random variable. The agent observes the objective probability p , chooses a stochastic perception of the probability p from the available set, forgets p , and forms a perception equal to the realization of the chosen random variable. The perception problem is to optimize over the perception strategies that

map from p to the set of stochastic perceptions. We assume that the menu of available stochastic perceptions is ordered with respect to first-order stochastic dominance and that the agent has more control in the high state than in the low one.

We compare solutions of two perception problems, both of which share the same menu of the stochastic perceptions but differ in the loss functions. In the *precision maximization* problem, the agent minimizes the mean-square error of her perception. In the *payoff maximization* problem, the agent minimizes the welfare loss. We use minimization of mean-square error as a benchmark because it is a widespread criterion in statistics and in information theory, even though it ignores asymmetries of the welfare losses from misperceptions. Comparison of the two solutions isolates the impact of the asymmetries on the payoff-maximizing perception. We show that, compared to the precision-maximizing benchmark, the payoff-maximizing strategy is biased upwards: the payoff-maximizing perception at each objective probability p stochastically dominates the precision-maximizing perception. Thus, an outsider observing the joint density of the objective and subjective probabilities and who conjectures that the agent is maximizing precision will unambiguously conclude that the agent is biased toward the high control state.

As a technical contribution, we provide a simple and intuitive formula that, for a given distribution of the utility functions, characterizes the loss from misperceiving the objective probability p as q . This characterization shows that the loss function can be expressed as an integral formula that depends only on the value function. The value function, for each probability p of the high state, specifies the payoff of the optimizing agent who has correct perception, averaged out across all the decision problems in her environment. The value function has a simple analytical form when the utility function has a Gaussian distribution, and hence we focus on the Gaussian payoffs in most of the paper. We obtain the asymmetry of the loss function and the bias of the optimal perception in settings in which (i) the payoff for each action consists of two additive terms, and (ii) one of the terms has a relatively higher weight in the high state than in the low state. We extend the results beyond the Gaussian payoff distributions in the last section of the paper.

Optimal information compression is studied in information theory through rate-distortion theory, as originated by Shannon (1948, 1959). This theory studies optimal communication via a noisy channel. One of the primitives is an exogenous map that specifies a loss to each input and output of the communication. A popular loss function is the square error. In our paper, the loss function is derived from the agent's environment and is equal to the

average welfare loss caused by the distortion.²

Once we establish the relevant loss function, we let the agent engage in optimal error management: she avoids the costlier types of the errors. Beyond information theory, the idea of error management has been studied in several scientific disciplines such as biology and psychology. In Haselton and Buss (2000), males' interpretation of females' sexual interest is biased toward false-positives because false-negatives have a larger fitness cost. Nettle (2004) provides a foundation for optimism by arguing for asymmetry of the optimistic and pessimistic error costs. See Johnson et al. (2013) for an interdisciplinary literature review and Alaoui and Penta (2016) for a recent axiomatization of the cost-benefit approach to error management within economics. We contribute to this literature a formal characterization of loss from misperception based on a statistical description of the agent's environment.

Sims (1998, 2003) has introduced information capacity constraints to economics. The subsequent literature on rational inattention studies the information acquisition and processing of an agent who, as in our model, is unsure about a payoff-relevant parameter but who, unlike our agent, knows her payoff function. In that setting (assuming signal cost is nondecreasing in Blackwell informativeness), the agent acquires only an action recommendation and no additional information beyond that needed for choice. In our model, the agent does not know her payoff function when she processes information and thus forms beliefs beyond those needed for the mere action choice in any given problem. The resulting optimal information structure can then be naturally interpreted as a perception of the payoff parameter.³

Our assumption that the perception strategy is optimized across many decision problems has appeared in the literature on the evolution of the utility functions. This literature studies performance of decision processes across a distribution of the fitness rewards to the available actions. Depending on the constraints assumed, the optimal decision criterion is either the expected utility maximization, as in Robson (2001), or its behavioral variants, as in Rayo and Becker (2007) and Netzer (2009). (See Robson and Samuelson (2010) for a literature review.) Most of this literature focuses on choice under certainty. Two exceptions who, like us, focus on the probability perception are Herold and Netzer (2010), who

²An appropriate definition of the loss function is a nontrivial and important part of the data compression problem. See the discussion of the loss function for compression of images in Wang et al. (2004).

³Within the rational inattention literature, Woodford (2012a,b) studies behavioral biases resulting from perception frictions; in these studies he assumes the mean square error loss criterion.

propose that biases in probability perception serve as a correction of another behavioral distortion; and Compte and Postlewaite (2012), who conclude that an agent with limited memory benefits from ignoring weakly informative signals.

We contribute to the branch of the bounded rationality literature that emphasizes limited memory. Mullainathan (2002) studies behavioral implications of exogenous imperfect memory usage. Like us, Dow (1991), Hirshleifer and Welch (2002) and Wilson (2014) study behavioral implications of an agent who optimizes her usage of limited memory, but they examine effects that are different from ours. In his survey, Lipman (1995) focuses on memory-related frictions.

This paper develops a model in the vein of Steiner and Stewart (2016) and focuses on new behavioral phenomena. A subtle but important difference in settings between the two papers is that Steiner and Stewart analyze perception of a probability that idiosyncratically affects the payoff to only one action, whereas this paper studies perception of a probability that affects payoffs to all actions. Because of this difference, the two papers arrive at opposite results. Steiner and Stewart derive the optimal overvaluation of small probabilities and offer an explanation for prospect-theory probability weighting. This paper, on the other hand, shows that the agent is overconfident about the precision of her belief and thus undervalues small probabilities attached to unlikely outcomes; it thus provides a microfoundation for the illusion of control and for overconfidence. In addition to the new application, the current paper makes a technical advance. The loss characterization based on the value function is novel and allows us to solve settings with large action sets.

Overconfidence, optimism, and other perception biases impact choices in personal, economic and political life.⁴ In Brunnermeier and Parker (2005), overconfidence bias arises because people derive hedonic payoffs from positive beliefs; in Bénabou and Tirole (2002) confidence helps to resolve an intrapersonal conflict caused by time-inconsistent preferences, while in Compte and Postlewaite (2004), confidence is an explicit production factor. Johnson and Fowler (2011) show that overconfidence provides a strategic advantage in a conflict. In our model, confidence has no intrinsic value and is a feature of the least costly feasible distribution of the perception errors.

⁴See empirical evidence for each of these domains in Puri and Robinson (2007), Hirshleifer (2001), and Ortoleva and Snowberg (2015).

2 Illustrative Example

The objective of this example is to illustrate our general model and methodology in a technically transparent setting. Section 3.1 presents three further applications that connect our model and results to well-known behavioral biases.

An agent has the choice as to whether to invest in a project or not. There are two states of the world, called high and low. The reward from investment is 1 in the high state and 0 in the low state, and the cost from investing is c in both states. The payoff from not investing is 0 in both states. An optimizing agent who attaches probability q to the high state chooses to invest if $q > c$ and not invest if $q < c$.

We allow for the possibility that the agent misperceives the state distribution, and study the loss incurred by such misperceptions. While the true probability of the high state is p , the probability perceived by the agent is q , and we let $L(p, q; c)$ be the expected loss from this misperception. If c does not lie between p and q , the agent makes the same decision under perception q as she would under p , so $L(p, q; c) = 0$. When c lies between p and q , the agent's choice under belief q differs from the optimal action under the true probability p , and $L(p, q; c) = |p - c|$ is the absolute difference between the expected profit $p - c$ from investing and the payoff 0 from not investing.

We are interested in situations in which errorless perception is impossible, so that the agent faces an error management problem. We assume that the error management takes place while the agent has incomplete information about her incentives; in this particular example, she manages errors while uncertain about the cost parameter c . This is relevant if the agent needs to assess the probability of the high state before the uncertainty about c is resolved, or if she cannot adjust her perception of the high state to different investment costs.

If c has density f over \mathbb{R}_+ , then the expected loss from misperceiving p as q is:

$$L(p, q) := \mathbb{E}_c L(p, q; c) = \int_0^\infty f(c)L(p, q; c)dc = \int_p^q f(c)(c - p)dc. \quad (1)$$

Below, we consider an agent who optimizes the distribution of perceived probabilities and study how her second-best perception strategy depends on the loss function. We focus on the case in which the loss function exhibits an asymmetry: an undervaluation of the

true probability of p is costlier than overvaluation of the same magnitude,

$$L(p, p - \sigma) > L(p, p + \sigma), \quad (2)$$

for all σ small enough so that $p - \sigma, p + \sigma \in [0, 1]$. We show that, when such asymmetry arises, the second-best perception strategy is biased upwards in the sense that $E[q | p] > p$. In this example, a sufficient condition for the asymmetry in (2) is that the density f is decreasing.⁵ The intuition behind this result is as follows. Higher values of c correspond to indifferences between the two actions for higher beliefs that the state is high. When f is decreasing, ties between investing and not investing are less likely to arise for beliefs in the interval $(p, p + \sigma)$ than in the interval $(p - \sigma, p)$. In this case, the misperception $p + \sigma$ of the probability p is less likely to distort the choice of action than the misperception $p - \sigma$. The next paragraph describes how this intuition drives the main technical insight of the paper.

The paper's main technical contribution relates the loss function to the value function of the optimizing agent. Let $v(p; c)$ be the value of the optimizing agent who correctly perceives the probability of the high state to be p . It is 0 for $p \leq c$, and $v(p; c) = p - c$ for $p \geq c$. The kink of $v(\cdot; c)$ at $p = c$ from a slope of 0 to a slope of 1 corresponds to a tie between investing and not investing arising at the belief $p = c$. The value of the optimizing agent who holds the correct belief p , averaged out across the random cost c is:

$$v(p) := E_c v(p; c) = E_c (p - c)^+ = \int_{c=0}^p f(c)(p - c)dc.$$

Two consecutive derivations of this formula allow to see that $v''(p)$ is well defined and equals $f(p)$. In this example, $v''(p)$ is the likelihood that the value function $v(\cdot; c)$ has a kink at p —that is that a tie arises at p . For general settings, Proposition 2 shows that $v''(p)$ is the expected intensity of a kink of $v(\cdot; c)$ at p —the likelihood of the tie multiplied by the average change of the slope of $v(\cdot; c)$ at p . Since in this setting all kinks have intensity 1, this equals the likelihood of the kink in the current example. We can thus rewrite (1) as

$$L(p, q) = \int_p^q v''(p^*)(p^* - p)dp^*.$$

⁵Formally, using a change of variable $c \mapsto 2p - c$ in the second term, one obtains: $L(p, p + \sigma) - L(p, p - \sigma) = \int_p^{p+\sigma} (f(c) - f(2p - c))(c - p)dc < 0$.

This characterization of losses is remarkable in the sense that the distribution of underlying decision problems is entirely factored through the second derivative v'' of the value function. Our main technical contribution, Lemma 1, shows that this characterization holds for general distributions of decision problems. It is particularly useful in settings such as the Gaussian setting where $L(p, q)$ cannot be easily computed directly, but v can. More importantly, it implies that the behavior of v'' defines the direction of the second-best perception biases. Intuitively, higher values of v'' correspond to more intense and more frequent kinks of the underlying payoff functions, they thus entail more frequent and more painful mistakes.

Our formulation of the loss function in terms of value function allows to neatly characterize the direction of biases in a class of settings with intuitive interpretation. In section 4.2 we consider decision problems with arbitrary finite set of actions $a \in A$ and binary state space $\theta \in \{h, l\}$ in which the payoffs $u(a, \theta)$ from action a in state θ are normally distributed. We identify conditions on the payoffs distribution under which the choice of the decision maker has a stronger impact on her utility in the high state than in the low one. In this case we say that the decision maker has *more control* in the high state than in the low one. We then show that v'' decreases with the probability attached to the high state, and the perception is biased towards the high-control state.

3 A Model of Misperceptions

An agent faces a randomly drawn decision problem. She chooses an action $a \in A = \{1, \dots, n\}$, $n \geq 2$, and receives payoff $u(a, \theta)$ where $\theta \in \{l, h\}$ is a payoff state. The agent's perception of the probability distribution over θ is imperfect. Correspondingly, we distinguish between an objective probability p and a subjective probability q of the state h . The objective and subjective probabilities of the state l are $1 - p$ and $1 - q$ respectively. The payoff functions $u(a, \cdot)$ are extended from $\theta \in \{l, h\}$ to all beliefs in $[0, 1]$ through the expected utility formula $u(a, p) = pu(a, h) + (1 - p)u(a, l)$.⁶

We study a two-stage decision process in which a *perception stage* is followed by a *choice stage*. In the *choice stage*, the agent knows the realized payoff function u , mistakenly

⁶ $p = 1$ is thus identified to $\theta = h$ and $p = 0$ to $\theta = l$.

perceives the true probability p as q , chooses an action⁷

$$a_{q,u}^* = \arg \max_{a \in A} u(a, q)$$

that maximizes her perceived payoff, and receives the true payoff $u(a_{q,u}^*, p)$.⁸

The decision problem is uncertain in the preceding *perception stage*. Accordingly, $u = (u(a, \theta))_{a, \theta}$ and p are random variables with supports \mathbb{R}^{2n} and $[0, 1]$, respectively, both admitting atomless density; u and p are independent.⁹ We let

$$V(p, q) = E_u u(a_{q,u}^*, p)$$

denote the payoff of the optimizing agent who misperceives the objective probability p as q , averaged out across the draws of the payoff function u . Let the value function $v(p) = V(p, p)$ denote the expected payoff of the optimizing agent who perceives the objective probability p correctly, and let

$$L(p, q) = v(p) - V(p, q)$$

be the expected loss due to misperceiving the true probability p as q .

Our aim is to characterize how the distribution of the payoffs u translates into the loss function L . To start, the next section reviews three applications.

3.1 Applications to Behavioral Biases

Overconfidence, overprecision,¹⁰ and illusion of control are well documented behavioral phenomena. The aim of the next three examples is to show that these systematic biases all arise as consequences of error management when the agent has imperfect perception of the world. Our model thus provides a unifying framework that encompasses these phenomena.

Example 1. [overconfidence] In this application, the agent is uncertain about her talent, and the choices of highly talented individuals are assumed to have on average more impact

⁷Note that this model is equivalent to one in which the agent's payoff is written as a function $g(a, \theta, \omega)$ that depends on the agent's action a , the payoff state θ , and a state of nature ω . Our modelling choice makes clear the fact that states of nature consist of (at most) $A \times \Theta$ parameters.

⁸In our framework, the maximizer is unique with probability one. We therefore ignore cases with multiple optima.

⁹Independence ensures that the agent does not draw inference about θ from the observation of the payoffs u , and vice versa. We discuss this assumption in section 8.

¹⁰We follow here the terminology from Moore and Healy (2008), who refer to overvaluation of one's performance as "overconfidence" and overvaluation of the precision of one's belief as "overprecision".

on their well-being than the choices of the less talented.

The agent chooses one out of n projects. Project $a \in \{1, \dots, n\}$ yields a payoff $u(a, \theta) = \tau_\theta b_a - c_a$, where c_a is the participation cost, $\tau_\theta > 0$ is the agent's talent in state θ , and b_a is the talent-unadjusted benefit of participating in project a . At the choice stage, the agent observes the $2n$ parameters $(b_a, c_a)_a$, attributes probability q to the high talent state, and chooses the project that maximizes $u(a, q) = qu(a, h) + (1 - q)u(a, l)$. The parameters b_a and c_a are independently drawn from a normal distribution in between the perception and choice stages.

We assume that $\tau_h > \tau_l$.¹¹ Misperception of talent leads to misinformed choice and a loss for some realizations of the projects' parameters. Our results show that the loss function is asymmetric in favor of the upward errors. Once we introduce optimization over the error distributions, this asymmetry will imply that when some misperception is unavoidable, the optimal perception exhibits overconfidence.

Example 2. [overprecision] In this example, the agent bets on an event. The agent may be either well informed or less informed about the event. On average, the better informed type of the agent faces larger payoff differences among the available bets than the less informed one.

The agent chooses one out of n binary bets. Bet $a \in \{1, \dots, n\}$ pays reward r_γ^a in outcome $\gamma \in \{t, d\}$, and the probabilities of the outcomes t and d in state $\theta \in \{h, l\}$ are π_θ and $1 - \pi_\theta$, respectively. Thus, the value of bet a in state θ is $u(a, \theta) = \pi_\theta r_t^a + (1 - \pi_\theta) r_d^a$. The agent observes the $2n$ realized rewards $(r_\gamma^a)_{a, \gamma}$ at the choice stage. She is uncertain about the state θ , and thus faces compound lotteries. The agent attributes probability q to state h and chooses the project that maximizes $u(a, q) = qu(a, h) + (1 - q)u(a, l)$. All rewards r_γ^a are independently drawn from a normal distribution in-between the agent's perception formation and her choice of bet.

We assume that $\pi_h > \pi_l > 1/2$ to reflect the fact the agent is more informed about the outcome of the bet in the high state than in the low one.¹² Misperception of θ leads to a mistaken belief about the outcome of the bet and to a suboptimal choice. Our results show that being underconfident about one's own knowledge of the payoff-relevant outcome

¹¹The optimal perception bias would be the same if we would assume that (i) same participation benefits are independent of the agent's talent level and (ii) the participation cost is increasing with talent, reflecting higher opportunity costs for the highly talented agents.

¹²The result is identical if $\pi_h < \pi_l < 1/2$.

is relatively costlier than the opposite error and that the second-best perception strategy exhibits overprecision.

Example 3. [illusion of control] In this example, the two states differ in the predictability of the payoff consequences of the agent's choice. The payoffs are more predictable in the high state than in the low one, and the agent's choice has on average more impact on her expected utility in the high state than in the low one.

The agent's payoff for action $a \in \{1, \dots, n\}$ is $e_a + l_a$, where e_a is a predictable payoff term (capturing, for instance, the predictable benefit of effort) and l_a is a luck term. At the choice stage, the agent observes, for each a , the predictable term e_a and a noisy signal $x_a = l_a + \varepsilon_a$ about the luck term, but not l_a itself. The luck terms l_a are drawn from $\mathcal{N}(y, 1)$ and the errors ε_a from $\mathcal{N}(0, 1/\gamma_\theta)$, where precision γ_θ represents the predictability of luck in state $\theta \in \{l, h\}$. Thus, the expectation over the luck term l_a upon observing signal x_a in the state θ is $\frac{y + \gamma_\theta x_a}{1 + \gamma_\theta}$, and the expected payoff for action a in state θ is $u(a, \theta) = e_a + \frac{y + \gamma_\theta x_a}{1 + \gamma_\theta}$. The agent observes the $2n$ parameters $(e_a, x_a)_a$ at the choice stage, attributes probability q to state h , and chooses the action that maximizes $u(a, q) = qu(a, h) + (1 - q)u(a, l)$. The terms e_a are independently drawn from a normal distribution in-between perception formation and choice of action.

We assume $\gamma_h > \gamma_l$ to express that luck is less predictable in the low state than in the high state. Misperception of p distorts the weight that the agent assigns to the signal x_a , and thus it distorts choice. We show that, again, the loss function is asymmetric in favor of the upward errors and the second-best perception strategy overestimates predictability of the payoffs, and thus also overestimates the impact of the agent's choice.

Proposition 1. *In all three examples, the downward perception error is costlier than the upward error of the same magnitude:*

$$L(p, p - \sigma) > L(p, p + \sigma),$$

for any $\sigma \in (0, 1/2)$ and $p \in [\sigma, 1 - \sigma]$.

The proposition follows from a general result in Corollary 1.

3.2 Intuition

Misperception of probability p as q causes a loss only when the optimal actions under the two probabilities differ: $a_{p,u}^* \neq a_{q,u}^*$. In this case, there exists a probability s in-between p

and q such that at s , the agent is indifferent between $a_{p,u}^*$ and $a_{q,u}^*$. Thus, misperception of p as q is costly for those draws of the payoffs u under which a tie arises in-between p and q . In all three examples, the variance of the payoff $u(a, s)$ is increasing in the belief s . In the first example, this is because the high talent magnifies the variance of the benefit parameter b_a . In the second example, the variance of the bet value $\pi r_t^a + (1 - \pi)r_d^a$ is $\pi^2 + (1 - \pi)^2$, which increases with π on $(1/2, 1)$. Thus, the more the agent knows about the outcome of the bet, the larger the payoff differences between the bet values tend to be. In the third example, the payoff consequences of each action are more predictable in the high state than in the low one, and thus the expected payoff differences are typically larger in the high state.

Larger values of the variance of $u(a, s)$ correspond to lower likelihoods of a tie between two actions at probability s . Consider these two subjective perceptions: $q = p + \sigma$ and $q = p - \sigma$. Due to the decreasing payoff variance, a tie is less likely to arise at some s in $(p, p + \sigma)$ than for $s \in (p - \sigma, p)$. Therefore, the upward perception error leads to a loss less often than the downward error of the same magnitude. The next section completes the above argument by incorporating into the computation the size of the loss from misperception alongside the probability of the loss.

4 Loss Characterization

Recall that $v(p) = E_u \max_a u(a, p)$ is the value of the optimizing agent with the correct perception, averaged out across the draws of the payoffs. Our main technical contribution characterizes the loss function in terms of the value function.

Lemma 1. *Assume v is twice differentiable on $[0, 1]$. Then, the expected payoff loss due to misperceiving the probability p as q is*

$$L(p, q) = \int_p^q v''(s)(s - p)ds. \quad (3)$$

We illustrate Lemma 1 in Figure 1. The left graph depicts functions $u(a, p)$ for all actions a under a particular realization of the payoffs u . The agent with perception q chooses the action associated with the full line, whereas the optimal choice at the objective probability p is the action associated with the dotted line, which implies the loss $L(p, q; u)$. When we let $v(p; u) = \max_a u(a, p)$, the loss is the error of the linear approximation of the




Figure 1: Left graph: loss from misperceiving p as q for a realized utility function. Right graph: Ex ante expected loss from misperceiving p as q .

realized value function $v(p; u)$ when it is approximated at q . The right graph depicts the expected loss as the error of the linear approximation of the expected value function.

In general, v might not be differentiable at every p , let alone twice differentiable. This is the case, for instance, if v consists of a single decision problem with finite choice set. We show, however, that v'' is well defined when the distribution ρ of decision problems u admits a continuous density. Further, $v''(p)$ admits a natural interpretation as the expected intensity of a kink at p over all u 's.

For each realization u of the agent's utility function, the value function $v(\cdot; u)$ in the decision problem with utility function u is given by:

$$v(p; u) = \arg \max_a E_p u(a, \theta).$$

The mappings $v(\cdot; u)$ are piecewise linear, and exhibit points of non-differentiability at beliefs in which the agent is indifferent between the two first-best actions. However, the right $v'_+(p; u)$ derivative and left derivative $v'_-(p; u)$ are well defined for all $0 < p < 1$, and we let $k(p; u) = v'_+(p; u) - v'_-(p; u)$ denote the intensity of the kink at p . If u has density ρ over $\mathbb{R}^{\Theta \times A}$, and if $U(p)$ denotes the set of parameters u such that $v(\cdot; u)$ has a kink at p , then the expected intensity of a kink at p is:

$$\int_{U(p)} k(p; u) d\rho(u).$$

The next proposition characterizes v'' as the expected intensity of a kink at p .

Proposition 2. *If the parameters $u(a, \theta)$ admit a continuous density over $\mathbb{R}^{A \times \theta}$, then v is twice differentiable on $[0, 1]$ and*

$$v''(p) = \int_{U(p)} k(p; u) d\rho(u).$$

4.1 The role of v''

The second derivative of the value function plays a central role in our analysis. We offer intuitive explanations of the role of v'' in our results.

Consider an agent with a prior belief p who receives a weakly informative signal upon which she updates her belief to π in a neighborhood of p . When the Bayesian constraint $E \pi = p$ and the second-order Taylor expansion are used, the value of the received information is

$$E v(\pi) - v(p) \approx v''(p) \text{Var}(\pi)/2, \quad (4)$$

where the information value results from the action adjustment to the updated belief.

We point out a geometric representation of the loss from a small misperception that highlights the fact that the loss arises due to distortion of choice. Notice in the left graph of Figure 1 that when the realized value function $v(\cdot; u)$ is linear in the neighborhood of p , the optimal action is identical for all the probabilities near p ; hence, a small misperception is inconsequential. A small misperception of the true probability p leads to a loss only if the realized value function exhibits a kink in a neighborhood of p where the kink represents an action reversal.

Now consider potentially large perception errors. $L(p, q; u)$ is positive when the optimal choices at p and q differ, and hence the loss is positive when the agent is indifferent between the first-best and the second-best action at some s between p and q . Figure 1 illustrates a situation in which such indifference arises at two values s and s' between p and q . The figure shows that, in this case, the loss $L(p, q; u)$ is the sum of two parts, L_s and $L_{s'}$, where each of these loss contributions is identified with a tie arising at probabilities s and s' . The ex ante expected loss $L(p, q)$ can thus be expressed as an integral $\int_p^q w(s)(s - p) ds$, where the weight $w(s)$ captures the likelihood of a tie at s and the intensity of the kink of $u(a, \cdot)$ at s . According to Proposition 2, this expected intensity $w(s)$ is simply $v''(s)$.

4.2 Gaussian Payoff Distributions and a Ranking of Control

We show here that the value function has a simple analytical form and the loss characterization becomes tractable when the payoff distribution is Gaussian. We address general distributions in section 7.

We assume in this section that the draws of the payoff values are independent across the actions a ; however, for each a , we allow for a correlation of the payoffs $u(a, l)$ and $u(a, h)$ across the two states. The payoff pair $(u(a, l), u(a, h))$ is drawn from a bivariate normal distribution $\mathcal{N}(\mu, \Sigma)$, where $\mu = (\mu_l, \mu_h)$ is the vector of means and $\Sigma = (\Sigma_{\theta, \theta'})_{\theta, \theta'}$ is the covariance matrix. The covariance allows us to accommodate applications in which the payoffs in the two states are generated by related processes, making the misperceptions natural.

When the likelihood of the high state is p , the expected payoff to each action a is normally distributed with mean $(1-p)\mu_l + p\mu_h$ and variance $(1-p)^2\Sigma_{ll} + 2p(1-p)\Sigma_{lh} + p^2\Sigma_{hh}$, and the payoffs from all actions are independent. The agent observes the decision problem parameters and chooses the action with the highest expected payoff. It follows that the value function of the decision problem can be written as:

$$v(p) = (1-p)\mu_l + p\mu_h + c\sqrt{(1-p)^2\Sigma_{ll} + 2p(1-p)\Sigma_{lh} + p^2\Sigma_{hh}}, \quad (5)$$

where the constant $c > 0$ is the expectation of the maximum of n independent draws from the standard normal distribution.¹³

This explicit expression for the value function allows us to characterize the loss function in any Gaussian setting, and more particularly in Examples 1-3. The technical property unifying these examples is that the payoff function can be decomposed as

$$u(a, \theta) = u_a + \tau_\theta \eta_a, \quad (6)$$

where u_a and η_a are independent normal random variables, iid. across actions, and $\tau_h > \tau_l > 0$.¹⁴ When such a decomposition holds, we say that the agent has *more control* in the high state than in the low one.

¹³Equation 5 relies on the scale and translation invariance of the normal distribution which ensures that the expected value of the maximum from n independent draws from a normal distribution $\mathcal{N}(\mu, \sigma^2)$ is $\mu + c\sigma$.

¹⁴See the appendix for this decomposition.

Lemma 2. *If the agent has more control in the high state than in the low one and the payoffs are normally distributed, then $v''(p)$ is decreasing.*

The next corollary follows from Lemmas 1 and 2.

Corollary 1 (asymmetric losses). *If the agent has more control in the high state than in the low one, then a downward perception error is costlier than an upward error of the same magnitude: $L(p, p - \sigma) > L(p, p + \sigma)$, for any $\sigma \in (0, 1/2)$ and all $p \in [\sigma, 1 - \sigma]$.*

Overconfidence in Example 1, overprecision in Example 2, and illusion of control of Example 3 all follow from Corollary 1.

5 Second-Best Perceptions

The first-best perception strategy—maximizing the expected payoff—sets $q = p$ for all p . We assume that perception is constrained to be noisy and we study second-best perception strategies.

We consider three specifications of the perception problem that differ in the constraints on the design of the decision process. We view the first specification as the simplest setting that allows us to discuss the second-best perception. The second setting relaxes the agent’s naiveté and the third setting offers testable implications under general perception technologies.

5.1 A Naive Perception Problem

To model the noise, a perception strategy associates to each objective probability p a distribution over the subjective perceptions q . In this section, a strategy assigns a stimulus $m(p)$ to each objective probability $p \in [0, 1]$, and a stimulus m triggers a perception $q = m + \varepsilon$, where ε is a binary random variable taking values σ and $-\sigma$ with equal probabilities, where $\sigma \in (0, 1/2)$. Thus, a *perception strategy* is a map $m : [0, 1] \rightarrow [\sigma, 1 - \sigma]$.¹⁵

A particular example of a perception strategy is the unbiased perception $m(p) = p$ for $p \in [\sigma, 1 - \sigma]$. This strategy does not exhibit a systematic error since $E[q | p] = p$ for all

¹⁵The simple additive error specification is not well suited for modelling perception of probabilities near the boundaries of the probability interval. See section 5.3 for a general perception technology that allows for nonadditive errors and can naturally accommodate probabilities near 0 and 1. Here, we ensure that every perception q is in the $[0, 1]$ interval by restricting m to $[\sigma, 1 - \sigma]$.

$p \in [\sigma, 1 - \sigma]$. The unbiased strategy is a benchmark to which we compare the optimal perception strategy $m^*(p)$ that solves a *naive perception problem* for each p ,

$$m^*(p) \in \arg \min_{m \in [\sigma, 1 - \sigma]} \mathbb{E}_\varepsilon L(p, m + \varepsilon). \quad (7)$$

Note that (7) is equivalent to payoff maximization.

We summarize the two-stage decision process as follows. The agent observes the objective probability p and memorizes a probability $m(p)$ at the perception stage, where $m(p)$ is independent of the payoff function u . At the choice stage, the agent recalls probability $q = m(p) + \varepsilon$ (but does not recall p and ε), and chooses an optimal action $a_{q,u}^*$ under her perception q and payoff u .

Lemma 1 yields a simple first-order condition.

Proposition 3. *Let $q = m^*(p) + \varepsilon$ be the stochastic perception under the perception strategy solving the naive perception problem (7) when the true probability is p . If the optimal stimulus $m^*(p)$ is in $(\sigma, 1 - \sigma)$, then*

$$\mathbb{E}_\varepsilon [v''(q)(q - p)] = 0. \quad (8)$$

Consider a stimulus $m \in [p - \sigma, p + \sigma]$.¹⁶ A marginal increase in m decreases the perception error $|m + \varepsilon - p|$ when $\varepsilon = -\sigma$, and increases it when $\varepsilon = \sigma$. A marginal change of the perception q affects the agent's payoff only if it affects her choice—that is, when a tie arises in-between the first-best and the second-best action at q . The first-order condition thus equalizes across the two perception errors the likelihood of a tie at $q = m + \varepsilon$, adjusted for the intensity of the kink and multiplied by the size of the misperception $|q - p|$.

When $v''(p)$ is monotone, the direction of the bias can be determined without fully solving for the optimal perception.

Proposition 4. *If $v''(p)$ is decreasing, then the perception strategy that solves the naive perception problem (7) is biased upwards: $m^*(p) > p$ and $\mathbb{E}[q | p] > p$ for all $p \in (0, 1 - \sigma)$.*

The next result summarizes conditions on the decision problem that generate frequent overestimation of the probability of the high-control state.

Corollary 2 (optimal illusion of control). *If the agent has more control in the high state than in the low one and the payoffs are normally distributed, then the perception strategy*

¹⁶Stimulus $m > p + \sigma$ is dominated by $m = p + \sigma$, and $m < p - \sigma$ is dominated by $m = p - \sigma$.

that solves the naive perception problem (7) is biased upwards: $m^*(p) > p$ and $E[q | p] > p$ for all $p \in (0, 1 - \sigma)$.

5.2 A Sophisticated Perception Problem

The agent from the previous section is naive in that she does not fully utilize the information available to her at the choice stage. This section studies the sophisticated perception strategies of an agent who loses some information during the decision process, but who is then fully capable of utilizing the information retained until the choice.

Let $M = \{0, 1\}$ be a binary message space. At the perception stage, the agent observes p , and retains information about it according to a *perception function* $m(p)$, $m : [0, 1] \rightarrow M$. At the choice stage, the agent chooses an action knowing payoffs $u = (u(a, \theta))_{a, \theta}$ and remembering the message $m(p)$. Her behavior at this stage is therefore described by a *choice rule* $c : \mathbb{R}^{2n} \times M \rightarrow A$. The *sophisticated perception problem* is:

$$\max_{m(\cdot), c(\cdot, \cdot)} E_{u, p} u(c(u, m(p)), p). \quad (9)$$

The next lemma states that the optimal perception function $m(\cdot)$ generates a binary partition $\mathcal{I}(p^*) = \{[0, p^*), [p^*, 1]\}$ of the probability interval. The agent at the choice stage remembers that p belongs to $I \in \mathcal{I}(p^*)$, forms a Bayesian posterior $q = E[p | p \in I]$, and chooses the optimal action $a_{q, u}^*$ under the posterior belief q .¹⁷

Lemma 3. *There exists p^* such that $m(p)$ is an indicator function of one of the two intervals of the partition $\mathcal{I}(p^*)$ for almost all p ; and $c(u, m) = a_{q(m), u}^*$ for almost all u , where $q(m)$ is the posterior belief $E[p | m(p) = m]$.*

Let us denote the density of p by φ and assume that it has support on $[0, 1]$ and is symmetric around half: $\varphi(p) = \varphi(1 - p)$ for all $p \in [0, 1]$. We interpret the symmetry as a neutrality assumption abstracting from biases driven by an asymmetry of the prior φ .

Denote the two Bayesian posteriors by $\underline{q}(p^*) = E[p | p < p^*]$, and $\bar{q}(p^*) = E[p | p \geq p^*]$.

Proposition 5 (sophisticated perception). *If $v''(p)$ is decreasing, and $\varphi(p)$ is symmetric around $1/2$, then the optimal threshold probability p^* is less than $1/2$ and the agent attains the high posterior $\bar{q}(p^*)$ with a higher probability than the low posterior $\underline{q}(p^*)$.*

¹⁷The partition of the probability interval is reminiscent of the coarse reasoning model in Jehiel (2005). One distinction is that our partition is optimal (among the binary partitions).

As seen in the proof in the appendix, Proposition 5 holds under a weaker condition on $v''(\cdot)$, namely, when $v''(p) > v''(1-p)$ for all $p \in [0, 1/2)$. This condition is satisfied in any Gaussian setting in which $\text{Var } u(a, h) > \text{Var } u(a, l)$; this inequality allows for any covariance of $u(a, l)$ and $u(a, h)$, and thus relaxes the additive structure of the payoffs from equation (6).

Let us summarize the results for the sophisticated perception.

Corollary 3. *If the agent has more control in the high state than in the low state, the payoffs are normally distributed, and $\varphi(p)$ is symmetric around $1/2$, then $p^* < 1/2$ and the agent attains the high posterior $\bar{q}(p^*)$ with a higher probability than the low posterior $\underline{q}(p^*)$.*

To obtain some intuition for the last result, recall from (4) that $v''(p)$ can be interpreted as a local value of information. Decreasing $v''(p)$ implies that the local value of information is higher in the lower half of the probability interval than in the upper one. The optimal binary partition then discriminates between the probabilities in the lower half of the probability interval at the expense of not distinguishing among the probabilities in the upper half.

The result of this section is related to the work by Menezes et al. (1980) on preferences against increasing downside risk. The authors study preferences over wealth lotteries with a common mean, and variances that differ in the size of the downside risk. They show that agents with a positive third derivative of the utility function prefer a smaller downside risk. If we interpret $v(p)$ as a “utility” over p , then our agent has a negative third derivative and thus prefers the “downside risk”—her optimal partition leads to a frequent but modest upward update to $\bar{q}(p^*)$ and an infrequent but large update to $\underline{q}(p^*)$.¹⁸

5.3 Testable Implications under General Perception Technologies

In this section we draw testable implications of our model. Since the analytical form of subjects’ perception technology cannot be known with certainty, we generalize results on the perception biases to a large class of perception technologies. Following a *revealed preference* type of approach, we study biases revealed by the data and provide general conditions under which our model predicts the observation of an upwards bias.

We take the perspective of an outsider who observes, for each value of the objective probability p , the conditional density $\mu(q|p)$ of the perceived probabilities q formed by the

¹⁸We thank Collin Raymond for pointing this connection out to us.

agent. Such data can be elicited in a lab experiment in which the experimenter controls the objective probability p and elicits the distribution of the perceptions q from choice data in many repetitions of the experiment.¹⁹ We provide a definition of the perception bias that relies only on the observed conditional densities and makes no assumptions about the agent’s perception technology. We then extend the results from the previous sections and show that the optimal perception is biased toward the high state if $v''(p)$ is decreasing and if the perception technology satisfies a regularity condition.

The benchmark against which we define the perception bias is a statistician whose sole goal is to maximize the precision of her perception, where precision is measured by the inverse of the mean-square error. As the agent’s true objective is payoff maximization and not precision maximization, the agent in general exhibits a different perception from that of the statistician. Our goal is to contrast the agent’s perception strategies with that of the statistician since such comparison reflects the asymmetries of the loss function that are absent from the statistician’s objective.

The following setting generalizes the two previous perception technologies. The set of available random perceptions is an arbitrary set Q , where each \mathbf{q} is a random variable with support $[0, 1]$. The agent observes the objective probability p and chooses a random perception $\mathbf{q} \in Q$. For each p , the agent solves the *general perception problem*

$$\min_{\mathbf{q} \in Q} \mathbb{E} L(p, \mathbf{q}). \tag{10}$$

The set Q of feasible random perceptions depends on the details of the perception technology available to the agent. For instance, for the technology from section 5.1, the set Q contains binary lotteries which attain values $m + \sigma$ and $m - \sigma$ with equal probabilities. For the technology from section 5.2, the set Q consists of two random variables with degenerate distributions, one of which attains the value $\underline{q}(p^*)$ with certainty and the other of which attains the value $\bar{q}(p^*)$. As in section 5.1, the set Q can be exogenous, or, as in section 5.2, the set Q may be the outcome of an optimization process.

Recall that the outsider observes the conditional density $\mu(q|p)$ for each p . When we let $\mathbf{q}(p)$ be a random variable with density $\mu(q|p)$, the observed data reveal the agent’s perception strategy $\mathbf{q} : [0, 1] \rightarrow Q$. We now provide a definition of when the choice data reveal that the agent is biased (compared to a statistician). The definition relies on the outsider’s knowledge of the perception strategy $\mathbf{q}(\cdot)$ but does not require the outsider to

¹⁹See, for instance, Wu and Gonzalez (1996) on the elicitation of such data.

know the full set of the feasible perceptions Q beyond the observed choices by the agent.

The outsider asks whether the agent can increase precision of her perception of the objective probability p by a deviation to her perception strategy $\mathbf{q}(\hat{p})$ applied to $\hat{p} \neq p$. Among all perceptions $\{\mathbf{q}(\hat{p}), \hat{p} \in [0, 1]\}$ in the dataset, a statistician would choose only perceptions that maximize precision, hence only perceptions $\mathbf{q}(\hat{p})$, where \hat{p} belongs to the set

$$P^*(p) = \arg \min_{\hat{p} \in [0, 1]} E(\mathbf{q}(\hat{p}) - p)^2.$$

When $p \notin P^*(p)$ then indeed the agent can increase precision by a deviation from $\mathbf{q}(p)$ to $\mathbf{q}(\hat{p})$, $\hat{p} \in P^*(p)$. Thus, when $p \notin P^*(p)$, the outsider can conclude from the choice data that the agent is not maximizing the statistician's objective and we say that the agent is *revealed to be biased at p* . Unsurprisingly, when the actual objective is payoff maximization, the agent is typically biased. Our contribution is in the characterization of the direction of the bias. We say that the agent is revealed to be biased *upwards* at p if

$$\sup P^*(p) < p.$$

Under this condition, the statistician with belief p benefits from choosing random perceptions that would have been chosen by types of the agent that are lower than p .²⁰

The next result holds under a regularity condition that requires the set Q to be ordered with respect to the strict stochastic dominance order. Recall that a random variable with cdf. F_X strictly stochastically dominates another random variable with cdf. F_Y if $F_X(z) \leq F_Y(z)$ for every z , and the inequality is strict for some z .

Proposition 6. *Assume that the set Q of the feasible perceptions is ordered with respect to the strict stochastic dominance order and that $v''(p)$ is decreasing. If the agent is revealed to be biased at p , then the agent is revealed to be biased upwards at p , and the payoff maximizing perception $\mathbf{q}(p)$ stochastically dominates the precision maximizing perception $\mathbf{q}(\hat{p})$, $\hat{p} \in P^*(p)$.*

The last result generalizes both Proposition 4 for the naive perception problem and Proposition 5 for the sophisticated case. When the perception error ε is additive as in section 5.1, the precision-maximizing perception strategy is the unbiased one; it sets random

²⁰In an alternative approach, one could look at the set of objective beliefs \hat{p} of the statistician that could have lead to a certain observed choice $\mathbf{q}(p)$ of the agent. The agent is then biased upwards if the revealed values of \hat{p} are all *higher* than the agent's objective belief p . In our framework, both approaches are equivalent.

perception at p to $p + \varepsilon$. Proposition 4 establishes that the payoff maximizing random perception at p is $m(p) + \varepsilon$ with $m(p) > p$. The agent is thus revealed to be biased upwards because the statistician would benefit from a downward shift of the agent’s strategy.

Let us verify that the proposition holds for the sophisticated perception from section 5.2. We have found that, for a symmetric prior $\varphi(p)$, the payoff-maximizing perception threshold is $p^* < 1/2$. The symmetry of $\varphi(p)$ implies that $p^* < (\bar{q}(p^*) + \underline{q}(p^*))/2$. For any given pair of perception levels $\underline{q} < \bar{q}$, the precision-maximizing strategy partitions the probability interval to $[0, p^{**}]$, and $(p^{**}, 1]$, with $p^{**} = (\bar{q} + \underline{q})/2$. Thus, the two perception strategies coincide for $p \notin (p^*, p^{**})$ and differ for $p \in (p^*, p^{**})$. When $p \in (p^*, p^{**})$, the payoff-maximizing agent forms perception \bar{q} , while the precision-maximizing agent forms \underline{q} . Thus, the agent is again revealed to be upward-biased at all $p \in (p^*, p^{**})$.

6 Related Perception Biases

This section clarifies the relationship of the illusion of control to two additional perception biases.

6.1 Optimism

Psychologists classify the illusion of control together with optimism among so-called positive illusions (Taylor and Brown, 1988). Our model helps to clarify the relationship between the two biases. Positive belief biases are often rationalized by assuming that the agent derives felicity from holding favorable beliefs; for example, see Brunnermeier and Parker (2005), Caplin and Leahy (2001), and Köszegi (2006). In contrast, the belief distortions are purely instrumental in our model; they guide choice. This section discusses the conditions under which the instrumental bias of this paper leans in the direction of optimism.

First, note that the optimal perception strategy in any of the perception problems of section 5 is invariant to any action-independent modification of the payoffs: the solution to a perception problem under stochastic payoffs u is also optimal under stochastic payoffs \tilde{u} , where $\tilde{u}(a, \theta) = u(a, \theta) + c_\theta$ for all (a, θ) , and c_θ are action-independent constants. This is so because the optimal perception depends only on $v''(p)$ whereas the constants c_θ affect the value function only by an additive linear term.

When we consider the naive perception technology from section 5.1, it is natural to interpret the agent as *optimistic* when $v(1) > v(0)$ and when the agent’s perception is

biased upwards in the sense that $m^*(p) > p$ for all $p \in (\sigma, 1 - \sigma)$. Note, however, that the optimal perception bias is caused by the ranking of control across the states, rather than by optimism. Adding constants to payoffs in either the high or the low state affects whether $v(1) > v(0)$ or $v(1) < v(0)$ but does not affect v'' and hence does not affect the optimal perception. Thus, the illusion of control does not generally entail optimism.

However, the illusion of control generates optimism under the additional assumption that the average payoff in the high state is at least as large as in the low state.

Proposition 7 (optimism). *If the agent has more control in the high state than in the low one, her payoffs are normally distributed, and $E_u u(a, h) \geq E_u u(a, l)$, then the agent is optimistic.*

Consider the worst case in which the means of the payoffs at the low and the high state are the same. The higher payoff variance in the high state implies that the optimizing agent can achieve a higher value at h than at l . Thus, unless the payoffs in the high-control state are predictably low, the illusion of control generates optimism.

6.2 Probability Weighting in Prospect Theory

Two well-known stylized phenomena, overconfidence and the prospect-theory probability weighting, have opposite implications for the perception of small probabilities. Overconfidence about one’s own knowledge implies an undervaluation of the small probabilities attached to unlikely outcomes, whereas prospect-theory predicts an overvaluation of small probabilities. Our model predicts both effects in two closely related settings and thus may help to clarify the respective domains of the two effects.²¹

We compare two settings. One of them is Example 2 and the second is the setting of Steiner and Stewart (2016). To facilitate the comparison, we describe both settings in a unified way. In both cases, the agent evaluates binary lotteries. For $a \in \{1, \dots, n\}$, lottery $\ell_a = (r_t^a, \pi_\theta; r_d^a, 1 - \pi_\theta)$ pays reward r_γ^a in outcome $\gamma \in \{t, d\}$. The rewards r_γ^a are iid. from the standard normal density and are observed by the agent at the choice stage. The probability π_θ of the outcome t is state-dependent, and $1/2 < \pi_l < \pi_h$. (The result is the same if π_θ is below half in both states.)

²¹The two phenomena were discussed by the same authors within a short period of time. Kahneman and Tversky (1977) discuss overconfidence, and Kahneman and Tversky (1979) discuss the overvaluation of small probabilities.

1. *Overprecision setting (example 2)*: The agent chooses one from n lotteries. The lotteries differ in the rewards r_γ^a but share the same probability parameter π_θ ; $u(a, \theta) = \pi_\theta r_t^a + (1 - \pi_\theta) r_d^a$ for $a \in \{1, \dots, n\}$.
2. *Prospect-theory setting*: The agent chooses lottery ℓ_1 or a safe option that pays s with certainty; $a \in \{1, 2\}$, $u(1, \theta) = \pi_\theta r_t^1 + (1 - \pi_\theta) r_d^1$ and $u(2, \theta) = s$.

In both settings, the agents solve the naive perception problem (7).

In the overprecision setting, all the lotteries are bets on the same event. The uncertainty about π thus affects payoffs to all actions. In the prospect-theory setting, the uncertainty affects only one action but not the agent's outside option.

Let $\pi(q) = \pi_h q + \pi_l(1 - q)$ be the subjective probability that the agent assigns to the likely outcome t ; note that it is above one-half. We have shown in section 5.1 that the agent's perception q is biased upwards in the overprecision setting. Thus, the perceived probability of the unlikely outcome d is biased downward; $E[1 - \pi(q) | p] < 1 - \pi(p)$. When the agent chooses from several binary bets on the same event, she underestimates the small probabilities.

The effect is reversed in the prospect-theory setting.

Proposition 8 (Steiner and Stewart, 2016). *If the outside option is $s > 3^{1/4}$, then the perception strategy that solves the perception problem (7) in the prospect-theory setting overvalues small probabilities: $E[1 - \pi(q) | p] > 1 - \pi(p)$.*

Steiner and Stewart interpret s as the maximum over the agent's alternatives to the lottery ℓ_1 . When the agent has sufficiently many alternatives, s is much more valuable than a typical draw of the lottery, and thus $s > 3^{1/4}$ is satisfied. The agent overvalues small probabilities that idiosyncratically affect the value of one of her actions but do not affect the alternatives.

The difference in the optimal biases across the two settings is caused by the difference in the pattern of the ties. Abusing notation slightly, let $v(\pi)$ be the value of the optimizing agent who correctly perceives the probability π of the outcome t . Figure 2 depicts $v''(\pi)$ in both settings. In the overprecision setting, a tie between the first-best and the second-best action is likely when the lottery value $r_t^a \pi + r_d^a(1 - \pi)$ has a *low* variance, which happens when π is close to half; $v''(\pi)$ is single-peaked with its maximum at half for the overprecision setting. In the prospect-theory setting, a tie between the lottery ℓ_1 and the high outside option s is likely when the lottery value $r_t^1 \pi + r_d^1(1 - \pi)$ has a *high* variance, and hence $v''(\pi)$




Figure 2: Solid line: $v''(\pi)$ in the overprecision setting. Dashed line: $v''(\pi)$ in the prospect-theory setting ($s = 1.5$).

is U-shaped in the prospect-theory setting. As the probability π is bounded above half in both settings, v'' exhibits opposite monotonicity across the two cases. Consequently, the optimal perception strategies have the opposite biases.

Kominers et al. (2015) offer an alternative model of frictional information processing and they too derive an optimal overvaluation of small probabilities. In their example 2, an agent allocates her unit endowment to two assets, t and d . In our notation, the agent can choose any lottery $(r_t, \pi; r_d, 1 - \pi)$ such that $r_t, r_d \in [0, 1]$ and $r_t + r_d = 1$. The agent has a logarithmic utility function, and thus her expected payoff from a lottery $(r_t, \pi; r_d, 1 - \pi)$ is $\pi \log r_t + (1 - \pi) \log r_d$. The optimal allocation for a given π is $r_t = \pi$, $r_d = 1 - \pi$, and therefore the value function of the optimizing agent with correct probability perception is $v(\pi) = \pi \log \pi + (1 - \pi) \log(1 - \pi)$. It is easy to verify that v'' is U-shaped with the minimum at half. Since v'' has the same monotonicity properties as the value function arising in the prospect-theory setting from above, the second-best perception biases in the two settings have the same direction.

7 General Payoff Distributions

We assume here that the agent has more control in the high state than in the low one. For the sake of simplicity, we consider only binary action sets in this section. That is, $u(a, \theta) = u_a + \tau_\theta \eta_a$, $\tau_h > \tau_l > 0$, $a \in \{1, 2\}$. We dispense with the assumption of Gaussian payoffs and provide sufficient conditions on general payoff distributions for the monotonicity

of $v''(p)$.

Let us introduce the differences $\delta = u_2 - u_1$, and $\eta = \eta_2 - \eta_1$; we denote their densities by $f(\delta)$ and $\phi(\eta)$. We allow for dependencies between u_1 and u_2 and between η_1 and η_2 but continue to assume that (u_1, u_2) and (η_1, η_2) are independent.

The next lemma formalizes the connection between the likelihood of ties and the second derivative of the value function. Let $\tau(p) = \tau_h p + \tau_l(1 - p)$, $\Delta = \delta + \tau(p)\eta$ be the payoff difference between the two actions, and p^* be the random belief at which the agent is indifferent between the two actions; that is for each δ and η , p^* solves $\delta + \tau(p^*)\eta = 0$. For each η , the conditional density of p^* is $f(-\tau(p)\eta) \left| \frac{d\Delta}{dp} \right| = f(-\tau(p)\eta) |(\tau_h - \tau_l)\eta|$. Thus, the unconditional density of p^* is

$$\int_{-\infty}^{\infty} f(-\tau(p)\eta) |(\tau_h - \tau_l)\eta| \phi(\eta) d\eta.$$

The intensity of the kink of the value function is $|(\tau_h - \tau_l)\eta|$. The next expression for $v''(p)$ augments the last integral by the intensity of the kinks.

Lemma 4.

$$v''(p) = (\tau_h - \tau_l)^2 \int_{-\infty}^{\infty} \eta^2 \phi(\eta) f(-\tau(p)\eta) d\eta. \quad (11)$$

The last lemma allows us to establish the monotonicity of $v''(p)$ under a regularity condition without computing $v(p)$. Recall that a function is single-peaked with its maximum at x if it is increasing below x and decreasing above x .

Lemma 5. *If the agent has more control in the high state than in the low one and the density $f(\delta)$ is single-peaked with its maximum at 0, then $v''(p)$ is decreasing on $(0, 1)$.*

The difference $\delta = u_2 - u_1$ has a single-peaked density when u_1 and u_2 are drawn iid. from many common densities, including the normal or Pareto densities. Additionally, we provide a condition on the joint distribution of u_1 and u_2 that is sufficient for the single-peakedness of $f(\delta)$. Recall that a function $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ is supermodular if for any $x, y, x', y' \in \mathbb{R}^4$:

$$h(\max\{x, x'\}, \max\{y, y'\}) + h(\min\{x, x'\}, \min\{y, y'\}) \geq h(x, y) + h(x', y').$$

Let $g(u_1, u_2)$ denote the joint density of u_1 and u_2 . We say that g is symmetric if $g(u_1, u_2) = g(u_2, u_1)$ for every (u_1, u_2) .

Lemma 6. *If g is symmetric and supermodular, then $f(\delta)$ is single-peaked with maximum at 0.*

The supermodularity assumption is related to the definition of affiliated random variables (Milgrom and Weber, 1982). The assumption generalizes the notion of positive correlation between u_1 and u_2 and corresponds to situations that may arise when the two payoff terms u_1 and u_2 are generated by a common process.²²

The following result concludes this section by summarizing conditions under which overconfidence obtains.

Corollary 4. *If the agent has more control in the high state than in the low one and the joint density g is both symmetric and supermodular, then $v''(p)$ is decreasing and the solution of the naive perception problem (7) is biased upwards: $m^*(p) > p$ and $E[q | p] > p$ for all $p \in (\sigma, 1 - \sigma)$. If, additionally, $\varphi(p)$ is symmetric around $1/2$, then the threshold probability p^* that solves the sophisticated perception problem (9) is less than $1/2$.*

8 Discussion

The designer of the decision process and the agent from our model can be interpreted in different ways. The designer may stand for the biological or cultural evolution of information processing whereas the agent may represent the individual. Alternatively, the model may be applied to the problem of the design and organization of a firm subject to frictions in its internal communication.

The model may help to inform us of the relevant domains of common behavioral phenomena. Though there is considerable consensus on both overconfidence and prospect-theory probability weighting, these two phenomena make conflicting predictions for the perception of small probabilities. A joint microfoundation may help to delineate this aspect of the two phenomena.

Models that provide microfoundations for behavioral biases can contribute to normative discussions of debiasing. Our model suggests that a bias relative to the precision-maximizing perception may be optimal and thus does not justify an intervention on its own. On the other hand, the model may inform us of circumstances under which an intervention into the decision process may be beneficial. A standard evolutionary argument

²²Random variables are affiliated if the logarithm of their joint density is supermodular. Our assumption applies to the density instead of the log-density.

for interventions applies: the modern agent's environment may differ from the one she is adapted to. In our case, we have assumed that the utility function u and the probability p are independent random variables. This is reasonable if the decision problem is generated by a nonstrategic process, but the independence assumption is violated in a market environment in which the reward for and probability of an outcome of a financial bet are typically negatively correlated.

The model suggests that an intervention may be justified even if the agent's environment has not changed since adaptation took place. To see this, consider a stylized example in which the intervening outsider knows more about the agent's decision problem than the perception designer has known. Assuming that the choices made by a healthy individual have a larger impact on her well-being than the choices made by a sick individual in most of the decision problems, the model predicts that health perception will be biased toward the denial of sickness. The assumption, however, may fail for a finer classification of the diseases. If some of the diseases are treatable by an appropriate action whereas some are not treatable, then the individual suffering from a treatable disease has considerable control over her well-being; the second-best perception should exhibit some hypochondria in the domain of treatable diseases and denial in the domain of nontreatable diseases. If the perception of one's own health status has been designed under a coarse classification of the diseases, whereas the intervening outsider distinguishes more finely between treatable and nontreatable diseases, then the outsider may call for an intervention even if the distribution of the diseases has not changed. We plan to return to a systematic study of the normative implications of the model in a future work.

A Proofs

Proof of Lemma 1. Note that $v(p) = \max_q V(p, q) = V(p, p)$. Applying the envelope theorem to

$$V(p, q) = \mathbb{E}_u [(1 - p)u(a_{q,u}^*, l) + pu(a_{q,u}^*, h)]$$

gives

$$v'(p) = \left. \frac{\partial}{\partial p} V(p, q) \right|_{q=p} = \mathbb{E}_u [u(a_{p,u}^*, h) - u(a_{p,u}^*, l)]. \quad (12)$$

Using the linearity of the expected payoff with respect to probability and (12), we obtain

$$V(p, q) = v(q) - v'(q)(q - p),$$

and thus,

$$L(p, q) = v(p) - V(p, q) = v(p) - v(q) + v'(q)(q - p)$$

is the error of the linear approximation of the value function $v(p)$ when it is approximated at q . Expression (3) follows once we integrate by parts:

$$\int_p^q v''(s)(s - p)ds = \int_p^q v'(s)ds - [(v'(q) - v'(s))(s - p)]_p^q = L(p, q).$$

□

Proposition 1 follows the proof of Lemma 2.

Proof of Proposition 2. Fix two actions a, a' , let $U_{a,a'}(p)$ be the set of values u for which $v(p; u)$ has a kink at p due to a switch from optimal action a to optimal action a' (or the reverse). This is the range of values for which there is indifference between a, a' at p and all other actions are inferior. Then we can decompose $v'(p)$ as

$$v'(p) = v'(0) + \sum_{a,a'} \int_{q=0}^p \int_{U_{a,a'}(q)} k(u; q) d\rho(u) dq.$$

Hence v'' is well defined and

$$v''(p) = \sum_{a,a'} \int_{U_{a,a'}(p)} k(p; u) d\rho(u) = \int_{U(p)} k(p; u) d\rho(u).$$

and $v''(p)$ equals the expected intensity of a kink at p over all u 's. \square

Proof of Lemma 2. Let $u_a \sim \mathcal{N}(\mu, \omega^2)$ and $\eta_a \sim \mathcal{N}(\nu, \omega'^2)$. Then

$$u(a, p) \sim \mathcal{N}(\mu + \tau(p)\nu, \omega^2 + \tau^2(p)\omega'^2),$$

where $\tau(p) = (1-p)\tau_l + p\tau_h$. Thus, $v(p) = \mu + \tau(p)\nu + c\sigma(p)$ where $\sigma(p) = \sqrt{\omega^2 + \tau^2(p)\omega'^2}$ is the standard deviation of $u(a, p)$. Observing that the linear term does not affect $v''(p)$, we have that $v''(p) = c\sigma''(p)$. Using the equality $\sigma(p) = \sqrt{c_1 + c_2p + c_3p^2}$ for some constants c_1, c_2 , and c_3 , we obtain that

$$v''(p) = c \frac{d^2}{dp^2} \sqrt{c_1 + c_2p + c_3p^2} = c \frac{-c_2^2 + 4c_1c_3}{4} \sigma^{-3}(p).$$

Since the value function is the expectation over the upper envelopes of linear functions and thus convex and since $\sigma(p)$ is positive, it must be the case that $c \frac{-c_2^2 + 4c_1c_3}{4}$ is positive. By the assumption that $\tau_h > \tau_l$, the function $\tau(p)$ and hence also $\sigma(p)$ are increasing. Thus, $v''(p)$ is decreasing in p . \square

Proof of Proposition 1. Example 1: Identification of $-c_a$ with u_a and b_a with η_a , respectively, from (6) implies that the agent has more control in the high state than in the low one.

Example 2: Let $u_a = (r_t^a + r_d^a)/2$, and let $\eta_a = r_t^a - r_d^a$; notice that they are independent. Then, $u(a, \theta) = u_a + (\pi_\theta - 1/2)\eta_a$. The assumption that $\pi_h > \pi_l > 1/2$ implies that the agent has more control in the high state than in the low one.

Example 3: Assume that the term e_a is drawn from $\mathcal{N}(\mu, \omega^2)$. Note that $u(a, \theta) \sim \mathcal{N}\left(\mu + y, \omega^2 + \frac{\gamma_\theta^2(1+1/\gamma_\theta)}{(1+\gamma_\theta)^2}\right) = \mathcal{N}\left(\mu + y, \omega^2 + \frac{\gamma_\theta}{\gamma_\theta+1}\right)$. Thus $u(a, \theta)$ equals in distribution $u_a + \sqrt{\frac{\gamma_\theta}{\gamma_\theta+1}}\eta_a$, where $u_a \sim \mathcal{N}(\mu + y, \omega^2)$, $\eta_a \sim \mathcal{N}(0, 1)$, and u_a, η_a are independent. By the assumed monotonicity we have $\gamma_h > \gamma_l$, it follows that $\sqrt{\frac{\gamma_h}{\gamma_h+1}} > \sqrt{\frac{\gamma_l}{\gamma_l+1}}$, and the agent has more control in the high state than in the low one.

The result now follows for all three examples from Corollary 1. \square

Proof of Proposition 3. By Lemma 1,

$$\mathbb{E}_\varepsilon L(p, m + \varepsilon) = \mathbb{E}_\varepsilon \int_p^{m+\varepsilon} v''(s)(s - p) ds.$$

Since

$$\frac{d}{dm} \mathbb{E}_\varepsilon \int_p^{m+\varepsilon} v''(s)(s-p)ds = \mathbb{E}_\varepsilon [v''(m+\varepsilon)(m+\varepsilon-p)],$$

the first-order condition for the interior optimal m gives (8). \square

Proof of Proposition 4. By Lemma 1,

$$\frac{d}{dm} \mathbb{E}_\varepsilon L(p, m+\varepsilon) = \mathbb{E}_\varepsilon [v''(m+\varepsilon)(m+\varepsilon-p)].$$

Assume by contradiction that $m \leq p$. Since $v''(p)$ is decreasing and $m - \sigma - p < 0$, we have:

$$\frac{d}{dm} \mathbb{E}_\varepsilon L(p, m+\varepsilon) < \mathbb{E}_\varepsilon [v''(m+\sigma)(m+\varepsilon-p)] = v''(m+\sigma)(m-p) \leq 0.$$

Since the constraint $m \leq 1-\sigma$ is not binding, the established inequality, $\frac{d}{dm} \mathbb{E}_\varepsilon L(p, m+\varepsilon) < 0$, violates the necessary first-order condition for the loss minimization. \square

Proof of Lemma 3. For any given $m(\cdot)$, the payoff optimization implies that $c(u, m) = a_{q(m), u}^*$, where $q(m) = \mathbb{E}[p \mid m(p) = m]$. For any two posteriors $\underline{q} < \bar{q}$, the difference between the losses across the two posteriors is

$$L(p, q') - L(p, q) = \int_q^{q'} v''(s)(s-p)ds.$$

The last expression is linear in p , and therefore there exists $p^* \in [0, 1]$ such that the posterior \underline{q} is optimal below p^* and \bar{q} is optimal above p^* . \square

Proof of Proposition 5. We prove the statement of the proposition under the condition

$$v''(p) > v''(1-p) \text{ for all } p < 1/2. \quad (13)$$

The monotonicity of $v''(p)$ assumed in the proposition implies (13).

Recall that $\underline{q}(p^*) = \mathbb{E}[p \mid p < p^*]$ and $\bar{q}(p^*) = \mathbb{E}[p \mid p \geq p^*]$. Note that the symmetry of the density of p implies $\underline{q}(p^*) < 1/2 < \bar{q}(p^*)$ for any $p^* \in (0, 1)$. Define

$$q(p; p^*) = \begin{cases} \bar{q}(p^*) & \text{if } p \geq p^*, \\ \underline{q}(p^*) & \text{if } p < p^*. \end{cases}$$

$q(p, p^*)$ denotes the agent's posterior when the objective probability is p and the threshold probability is p^* .

We compare the payoff of the agent who uses the threshold p^* with the payoff of an agent who has only the prior information about p . By the symmetry of the prior density of p , this last agent attaches probability $\mathbb{E} p = 1/2$ to the high state and her expected payoff is:

$$v(1/2) = \mathbb{E}_p V(q(p; p^*), 1/2).$$

The expected payoff to the agent who observes whether $p < p^*$ or $p \geq p^*$ is $\mathbb{E}_p v(q(p; p^*))$ and the optimal partition maximizes this value over p^* . We rewrite this objective function as:

$$\begin{aligned} \mathbb{E}_p v(q(p; p^*)) &= v(1/2) + \mathbb{E}_p [v(q(p; p^*)) - v(1/2)] \\ &= v(1/2) + \mathbb{E}_p [v(q(p; p^*)) - V(q(p; p^*), 1/2)] \\ &= v(1/2) + \mathbb{E}_p L(q(p; p^*), 1/2). \end{aligned}$$

Notice that $\Pr(p < p^*) = \frac{\bar{q}(p^*) - 1/2}{\bar{q}(p^*) - \underline{q}(p^*)}$ and $\Pr(p \geq p^*) = \frac{1/2 - \underline{q}(p^*)}{\bar{q}(p^*) - \underline{q}(p^*)}$. Thus, the threshold p^* maximizes the objective

$$o(p^*) := \frac{\bar{q}(p^*) - 1/2}{\bar{q}(p^*) - \underline{q}(p^*)} L(\underline{q}(p^*), 1/2) + \frac{1/2 - \underline{q}(p^*)}{\bar{q}(p^*) - \underline{q}(p^*)} L(\bar{q}(p^*), 1/2). \quad (14)$$

Substituting (3) for $L(q, 1/2)$ into (14) gives

$$o(p^*) = \int_0^1 v''(s) h(s, p^*) ds,$$

where

$$h(s, p^*) = \begin{cases} 0 & \text{if } p < \underline{q}(p^*), \\ \frac{\bar{q}(p^*) - 1/2}{\bar{q}(p^*) - \underline{q}(p^*)} (s - \underline{q}(p^*)) & \text{if } \underline{q}(p^*) \leq p < 1/2, \\ \frac{1/2 - \underline{q}(p^*)}{\bar{q}(p^*) - \underline{q}(p^*)} (\bar{q}(p^*) - s) & \text{if } 1/2 \leq p < \bar{q}(p^*), \\ 0 & \text{if } p > \bar{q}(p^*). \end{cases} \quad (15)$$

The function $h(s, p^*)$ is continuous piecewise-linear in s , as depicted in figure 3.

Let us now prove that $o(p^*) > o(1 - p^*)$ when $p^* < 1/2$. For such p^* , the symmetry of




Figure 3: $p^* < 1/2$. Thick line: $h(s, p^*)$. Dashed line: $h(s, 1 - p^*) = h(1 - s, p^*)$. Note that $h(s, 1 - p^*) \leq h(s, p^*)$ for $s \leq 1/2$.

the density of p implies that $q(p^*) < 1 - \bar{q}(p^*)$. The definition (15) of h then implies

$$h(s, p^*) > h(1 - s, p^*) \text{ for all } p^* < 1/2, s < 1/2. \quad (16)$$

Note also that by the symmetry of $\varphi(p)$, it follows that:

$$h(s, p^*) = h(1 - s, 1 - p^*) \text{ for all } p^*, s \in (0, 1).$$

We now obtain

$$\begin{aligned} o(p^*) - o(1 - p^*) &= \int_0^1 v''(s) (h(s, p^*) - h(s, 1 - p^*)) ds \\ &= \int_0^1 v''(s) (h(s, p^*) - h(1 - s, p^*)) ds \\ &= \int_0^{1/2} (v''(s) - v''(1 - s)) (h(s, p^*) - h(1 - s, p^*)) ds. \end{aligned}$$

The first bracket in the last integral is positive by condition (13). When $p^* < 1/2$, the second bracket is positive by (16). We have thus established that $p^* > 1/2$ cannot be optimal. Hence, $p^* \leq 1/2$.

It remains for us to prove that $p^* \neq 1/2$. Let

$$w(p^*, q_l, q_h) = \int_0^{p^*} V(p, q_l) \varphi(p) dp + \int_{p^*}^1 V(p, q_h) \varphi(p) dp$$

stand for the expected payoff when the agent forms belief q_l for $p < p^*$ and belief q_h for $p \geq p^*$. The agent's objective is $E_p v(q(p; p^*)) = w(p^*, \underline{q}(p^*), \bar{q}(p^*))$. Notice that

$$\begin{aligned} \frac{d}{dp^*} E_p v(q(p; p^*)) &= w_1(p^*, \underline{q}(p^*), \bar{q}(p^*)) \\ &= \varphi(p^*) (V(p^*, \underline{q}(p^*)) - V(p^*, \bar{q}(p^*))) \\ &= \varphi(p^*) (V(p^*, \underline{q}(p^*)) - v(p^*) - V(p^*, \bar{q}(p^*)) + v(p^*)) \\ &= \varphi(p^*) (L(p^*, \bar{q}(p^*)) - L(p^*, \underline{q}(p^*))), \end{aligned}$$

where w_1 in the first line stands for the partial derivative with respect to the first argument. The first equality follows from the optimality of the Bayesian posteriors and from the Envelope Theorem. Observe that, by the symmetry of φ , it follows that $\underline{q}(1/2) = 1 - \bar{q}(1/2)$. Thus,

$$\begin{aligned} \left. \frac{d}{dp^*} E_p v(q(p; p^*)) \right|_{p^*=1/2} &= \varphi(1/2) (L(1/2, 1 - \underline{q}(1/2)) - L(1/2, \underline{q}(1/2))) \\ &= \varphi(1/2) \left(\int_{1/2}^{1-\underline{q}(1/2)} v''(s)(s - 1/2) ds - \int_{1/2}^{\underline{q}(1/2)} v''(s)(s - 1/2) ds \right) \\ &= \varphi(1/2) \left(\int_{1/2}^{1-\underline{q}(1/2)} (v''(s) - v''(1-s))(s - 1/2) ds \right), \end{aligned}$$

and the last integral is negative by condition (13). Thus, $p^* = 1/2$ cannot satisfy the first-order condition necessary for the optimum. \square

Proof of Proposition 6. We compare the solutions to two problems. Let $Q^* = \{\mathbf{q}(p) : p \in [0, 1]\}$ be the set of random perceptions used by the agent. The payoff-maximizing perception solves

$$\mathbf{q}(p) \in \arg \max_{\mathbf{q} \in Q^*} E[-L(p, \mathbf{q})],$$

and the precision-maximizing perception solves

$$\hat{\mathbf{q}}(p) \in \arg \max_{\mathbf{q} \in Q^*} \mathbf{E} \left[-\hat{L}(p, \mathbf{q}) \right],$$

where

$$\hat{L}(p, q) = \frac{v''(p)}{2}(q-p)^2 = \int_p^q v''(p)(s-p)ds.$$

The expression $\hat{L}(p, q)$ differs from the square of the perception error only by the multiplicative factor $v''(p)/2$, which does not affect the solution of the precision optimization.

First, note that the objective function $-L(p, q)$ has increasing differences in (p, q) . Indeed, by Lemma 1,

$$-L(p, q') + L(p, q) = \int_q^{q'} (p-s)v''(s)ds,$$

which increases in p when $q' > q$. Thus, the objective $\mathbf{E}[-L(p, \mathbf{q})]$ has increasing differences (with \mathbf{q} ordered by the stochastic dominance order). Therefore, by Theorem 1 in Van Zandt (2002), the payoff-maximizing perception $\mathbf{q}(p)$ is non-decreasing in p with respect to the stochastic dominance order.

Further, note that the objective $-L(p, q)$ has larger payoff differences than the objective $-\hat{L}(p, q)$. Indeed, using Lemma 1,

$$-L(p, q') + L(p, q) - \left(-\hat{L}(p, q') + \hat{L}(p, q) \right) = \int_q^{q'} (s-p)(v''(p) - v''(s)) ds.$$

The last expression is positive if $q' > q$ because v'' is decreasing. Thus, the objective $\mathbf{E}[-L(p, \mathbf{q})]$ has larger payoff differences than $\mathbf{E}[-\hat{L}(p, \mathbf{q})]$. This means that, using \succ to denote the first-order stochastic dominance order, for any p and $\mathbf{q}' \succ \mathbf{q}$,

$$\mathbf{E}[-L(p, \mathbf{q}')] - \mathbf{E}[-L(p, \mathbf{q})] > \mathbf{E}[-\hat{L}(p, \mathbf{q}')] - \mathbf{E}[-\hat{L}(p, \mathbf{q})].$$

Therefore, by Theorem 1 in Van Zandt (2002), $\mathbf{q}(p) \succeq \hat{\mathbf{q}}(p)$ for all p and for all solutions $\mathbf{q}(p)$ and $\hat{\mathbf{q}}(p)$ of the two problems. Since, because the agent is revealed to be biased at p , $\mathbf{q}(p)$ does not solve the precision optimization problem, it must be that $\mathbf{q}(p) \succ \hat{\mathbf{q}}(p)$. Since $\hat{\mathbf{q}}(p) \in Q^*$, there exists $\hat{p} \in [0, 1]$ such that $\hat{\mathbf{q}}(p) = \mathbf{q}(\hat{p})$ and because $\mathbf{q}(p)$ is non decreasing, $\hat{p} < p$ for all such \hat{p} . \square

Proof of Proposition 7. It suffices to prove the proposition for the case $\mathbb{E}_u[u(a, h)] = \mathbb{E}_u[u(a, l)]$, because $v(1)$ increases with an upward shift of the distribution of $u(a, h)$. More control in the high state implies that $\text{Var}_u[u(a, h)] > \text{Var}_u[u(a, l)]$. Together with $\mathbb{E}_u[u(a, h)] = \mathbb{E}_u[u(a, l)]$ and normality, this implies that the payoffs in the high state are a mean-preserving spread of the payoffs in the low state. By convexity of the max function,

$$v(1) = \mathbb{E} \max_a \{u(a, h)\} > \mathbb{E} \max_a \{u(a, l)\} = v(0).$$

By Corollary 2, the optimal perception is biased upwards. Hence, the bias is optimistic. \square

Proposition 8 was proved in Steiner and Stewart (2016); see their Theorem 1. Here, we briefly sketch the proof of the proposition using the analytical approach of this paper. Let $v(p) = \mathbb{E}_{r_t, r_d} \max\{\pi(p)r_t + (1 - \pi(p))r_d, s\}$. By Proposition 4 (reversing the labels of the states), it suffices to prove that $v''(p)$ is increasing. Since $\pi(p)$ is a linear mapping that attains values in $(1/2, 1)$, it suffices to prove that $\tilde{v}''(\pi)$ is increasing on $(1/2, 1)$, where $\tilde{v}(\pi) = \mathbb{E}_{r_t, r_d} \max\{\pi r_t + (1 - \pi)r_d, s\}$. Observe that

$$\tilde{v}'(\pi) = \mathbb{E}_{r_t, r_d} [(r_t - r_d) \mathbf{1}_{\pi r_t + (1 - \pi)r_d \geq s}].$$

where $\mathbf{1}_{\pi r_t + (1 - \pi)r_d \geq s}$ is 1 if $\pi r_t + (1 - \pi)r_d \geq s$, and zero otherwise. One can verify that $\tilde{v}''(\pi) = \lambda(\pi)w(\pi)$, where $w(\pi)$ is the likelihood of a tie at $\pi r_t + (1 - \pi)r_d = s$, and $\lambda(\pi) = \mathbb{E}[(r_t - r_d)^2 \mid \pi r_t + (1 - \pi)r_d = s]$. Steiner and Stewart (2016) provide explicit forms of $w(\pi)$ and $\lambda(\pi)$ and prove that if $s > 3^{1/4}$ then $\lambda(\pi)w(\pi)$ is increasing on $(1/2, 1)$; see their Lemma 2.

Proof of Lemma 4. Using insight (12) from the beginning of the proof of Lemma 1, we obtain:

$$\begin{aligned} v'(p) &= \mathbb{E}_u [u(a_{p,u}^*, h) - u(a_{p,u}^*, l)] \\ &= \mathbb{E}_u [u_{a_{p,u}^*} + \tau_h \eta_{a_{p,u}^*} - u_{a_{p,u}^*} - \tau_l \eta_{a_{p,u}^*}] \\ &= (\tau_h - \tau_l) \mathbb{E}_u \eta_{a_{p,u}^*}. \end{aligned} \tag{17}$$

Recall that $\tau(p) = p\tau_h + (1 - p)\tau_l$. The agent chooses action 1 if $-\tau(p)\eta > \delta$, and chooses

action 2 if $-\tau(p)\eta < \delta$. Thus,

$$\begin{aligned} E_u \eta_{a_{p,u}^*} &= \int_{\{(\eta_1, \eta_2, \delta): -\tau(p)\eta > \delta\}} \eta_1 f(\delta) \phi(\eta) d\delta d\eta_1 d\eta_2 + \int_{\{(\eta_1, \eta_2, \delta): -\tau(p)\eta < \delta\}} \eta_2 f(\delta) \phi(\eta) d\delta d\eta_1 d\eta_2 \\ &= \int_{\mathbb{R}^2} \phi(\eta) \left(\int_{-\infty}^{-\tau(p)\eta} \eta_1 f(\delta) d\delta + \int_{-\tau(p)\eta}^{\infty} \eta_2 f(\delta) d\delta \right) d\eta_1 d\eta_2, \end{aligned}$$

and

$$\begin{aligned} \frac{d}{dp} E_u \eta_{a_{p,u}^*} &= \int_{\mathbb{R}^2} \phi(\eta) \left(-\eta_1 f(-\tau(p)\eta) (\tau_h - \tau_l) \eta + \eta_2 f(-\tau(p)\eta) (\tau_h - \tau_l) \eta \right) d\eta_1 d\eta_2 \\ &= (\tau_h - \tau_l) \int_{-\infty}^{\infty} \phi(\eta) \eta^2 f(-\tau(p)\eta) d\eta. \end{aligned}$$

Substitution of the last line into (17) gives (11). \square

Proof of Lemma 5. The result follows from (11): it follows from $\tau_h > \tau_l$ that $\tau(p)$ is increasing in p , which implies that $f(-\tau(p)\eta)$ is decreasing in p for all p and η . \square

Proof of Lemma 6. The symmetry of g implies $f(\delta) = f(-\delta)$ for all δ . Thus, for arbitrary $t \in \mathbb{R}$,

$$f(\delta) = \int_{-\infty}^{\infty} g(x+t, x+t+|\delta|) dx = \int_{-\infty}^{\infty} g(x+t, x+t-|\delta|) dx. \quad (18)$$

Consider δ, δ' such that $|\delta'| > |\delta|$, and let $z = |\delta'| - |\delta| > 0$. Using (18), we get:

$$2(f(\delta) - f(\delta')) = \int_{-\infty}^{\infty} g(u, u+|\delta|) + g(u-z, u-z-|\delta|) - g(u, u-|\delta'|) - g(u-z, u-z+|\delta'|) du.$$

Since

$$u+|\delta| = u-z+|\delta'| \text{ and } u-|\delta'| = u-z-|\delta|,$$

the supermodularity of g implies that the last integrand is positive. \square

References

- Alaoui, L. and A. Penta (2016). Cost-benefit analysis in reasoning. mimeo.
- Bénabou, R. and J. Tirole (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics*, 871–915.

- Brunnermeier, M. K. and J. A. Parker (2005). Optimal expectations. *American Economic Review* 95(3), 1092–1118.
- Caplin, A. and J. Leahy (2001). Psychological expected utility theory and anticipatory feelings. *Quarterly Journal of Economics*, 55–79.
- Compte, O. and A. Postlewaite (2004). Confidence-enhanced performance. *American Economic Review*, 1536–1557.
- Compte, O. and A. Postlewaite (2012). Belief formation. PIER Working Paper 12-027, University of Pennsylvania.
- Dow, J. (1991). Search decisions with limited memory. *The Review of Economic Studies* 58(1), 1–14.
- Haselton, M. G. and D. M. Buss (2000). Error management theory: A new perspective on biases in cross-sex mind reading. *Journal of Personality and Social Psychology* 78(1), 81–91.
- Herold, F. and N. Netzer (2010). Probability weighting as evolutionary second-best. University of Zurich Socioeconomic Institute Working Paper 1005.
- Hirshleifer, D. (2001). Investor psychology and asset pricing. *The Journal of Finance* 56(4), 1533–1597.
- Hirshleifer, D. and I. Welch (2002). An economic approach to the psychology of change: Amnesia, inertia, and impulsiveness. *Journal of Economics & Management Strategy* 11(3), 379–421.
- Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory* (123), 81–104.
- Johnson, D. D., D. T. Blumstein, J. H. Fowler, and M. G. Haselton (2013). The evolution of error: Error management, cognitive constraints, and adaptive decision-making biases. *Trends in Ecology & Evolution* 28(8), 474–481.
- Johnson, D. D. and J. H. Fowler (2011). The evolution of overconfidence. *Nature* 477(7364), 317–320.

- Kahneman, D. and A. Tversky (1977). Intuitive prediction: Biases and corrective procedures. Technical report, DTIC Document.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291.
- Kominers, S., X. Mu, and A. Peysakhovich (2015). Paying (for) attention: The impact of information processing costs on bayesian inference. Technical report, Harvard University.
- Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association* 4(4), 673–707.
- Langer, E. J. (1975). The illusion of control. *Journal of personality and social psychology* 32(2), 311–328.
- Lipman, B. L. (1995). Information processing and bounded rationality: A survey. *Canadian Journal of Economics*, 42–67.
- Menezes, C., C. Geiss, and J. Tressler (1980). Increasing downside risk. *American Economic Review* 70(5), 921–932.
- Milgrom, P. R. and R. J. Weber (1982). A theory of auctions and competitive bidding. *Econometrica*, 1089–1122.
- Moore, D. A. and P. J. Healy (2008). The trouble with overconfidence. *Psychological Review* 115(2), 502–517.
- Mullainathan, S. (2002). A memory-based model of bounded rationality. *Quarterly Journal of Economics*, 735–774.
- Nettle, D. (2004). Adaptive illusions: optimism, control and human rationality. In D. Evans and P. Cruse (Eds.), *Emotion, Evolution and Rationality*, pp. 193–208. Oxford University Press.
- Netzer, N. (2009). Evolution of time preferences and attitudes toward risk. *American Economic Review* 99(3), 937–955.
- Ortoleva, P. and E. Snowberg (2015). Overconfidence in political behavior. *American Economic Review* 105(2), 504–535.

- Puri, M. and D. T. Robinson (2007). Optimism and economic choice. *Journal of Financial Economics* 86(1), 71–99.
- Rayo, L. and G. S. Becker (2007). Evolutionary efficiency and happiness. *Journal of Political Economy* 115(2), 302–337.
- Robson, A. J. (2001). Why would nature give individuals utility functions? *Journal of Political Economy* 109(4), 900–914.
- Robson, A. J. and L. Samuelson (2010). The evolutionary foundations of preferences. In *Handbook of social economics*, Volume 1, pp. 221–310. Elsevier, Amsterdam: North Holland.
- Shannon, C. (1959). Coding theorems for a discrete source with a fidelity criterion. *Institute of Radio Engineers National Convention Record part 4* 7(142-163), 142–163.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal* 27, 3–55.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 99–118.
- Sims, C. A. (1998). Stickiness. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 49, pp. 317–356. Elsevier.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–690.
- Steiner, J. and C. Stewart (2016). Perceiving prospects properly. *American Economic Review* 106(7), 1601–31.
- Taylor, S. E. and J. D. Brown (1988). Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin* 103(2), 193.
- Van Zandt, T. (2002). An introduction to monotone comparative statics. Insead, Fontainebleau.
- Wang, Z., A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13(4), 600–612.

- Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica* 82(6), 2257–2294.
- Woodford, M. (2012a). Inattentive valuation and reference-dependent choice. Columbia University, New York.
- Woodford, M. (2012b). Prospect theory as efficient perceptual distortion. *American Economic Review* 102(3), 41–46.
- Wu, G. and R. Gonzalez (1996). Curvature of the probability weighting function. *Management Science* 42(12), 1676–1690.