Naive Calibration

Yair Antler* and Benjamin Bachi‡

May 15, 2023

Abstract

We develop a model of non-Bayesian decision-making in which an agent obtains a signal about a relevant economic fundamental and subsequently takes an action. To interpret the signal, the agent calibrates a simple prediction rule based on a dataset that consists of previous signals and state realizations. Her subsequent action affects the probability with which the current signal and the corresponding state realization will be observed and recorded in the dataset that will be used in future decisions. We show that this procedure converges to a steady state and that it results in a seemingly pessimistic behavior that is exacerbated by feedback loops. We apply our model to project selection problems and second-price IPV auctions.

1 Introduction

Imagine that you are a manager of a company that has to make an important decision whose outcome depends on an economic variable, and that you have a noisy estimate of the variable's value. You look in the company's records and find that in similar decision problems faced by your predecessors, on average, the estimates at their disposal were

*Coller School of Management, Tel Aviv University. *E-mail:* yair.an@gmail.com.

†Department of Economics, University of Haifa. *E-mail:* bbachi@econ.haifa.ac.il

‡We thank Daniel Bird, Duarte Gonçalves, Philippe Jehiel, Botond Kőszegi, Dotan Persitz, Ariel Rubinstein, and Rani Spiegler for helpful comments and suggestions. We also thank seminar audiences at BRIC 2022, ESEM-EEA 2022, Hebrew University, NASM 2022, NYU Abu Dhabi, Tel Aviv University, UAB, UCL, and the University of Haifa for useful comments. Antler gratefully acknowledges financial support from the Coller Foundation and the Henry Crown Institute of Business Research in Israel.

10% higher than the actual outcomes. How would you interpret the current estimate in light of this finding?

A natural way to account for the apparent bias in the company's records is to adjust the estimate and discount it by 10%. This type of adjustment is called *reference class forecasting* and its use is advocated by researchers, government agencies, and professional associations in order to account for optimism bias and to prevent cost overruns in construction projects (Kahneman and Lovallo, 1993; HM Treasury, 2003a,b; AACE International, 2012; Flyvbjerg, 2008). In particular, the UK Department for Transport has employed it in the appraisal process of the cost and the construction duration of large transportation projects (Flyvbjerg et al., 2004; UK Department for Transport and Oxford Global Projects, 2020). Similar procedures are used in the oil and gas extraction industries (Nesvold and Bratvold, 2022). The need for such adjustments stems from the fact that oil extraction projects are often selected based on their estimated reserves and, in ex post audits, these reserves are typically lower than initially estimated (Brashear et al., 2001; Chen and Dyer, 2009).

One problem with this approach is that systematic differences between estimates and ex post outcomes in the dataset may emerge even when the estimates are unbiased if the data suffers from selection bias (e.g., the company's records may depend on one's predecessors' decisions or the organizational memory). It is well known that when failing to properly account for selection bias, one may end up reaching an erroneous conclusion. But what happens when the person who reaches the erroneous conclusion is not an outside observer but rather a decision maker whose actions affect the selection bias in the data she is using? To address this question, we develop a model in which boundedly rational DMs try to account for a selection bias they inadvertently contributed to themselves.

To illustrate our conceptual framework and some of our findings, consider the next example, which is inspired by Jehiel's (2018) model of investment decisions.

Example 1 *A risk-neutral entrepreneur decides whether to implement risky projects or not based on their estimated returns. Implementing a project costs* 0.4 *and yields a revenue of* $\theta \sim U[0,1]$. *The estimated return,* s, *equals* θ *with probability* $p < 1$. *Otherwise,* s *is independently drawn from* $U[0,1]$.

As a benchmark, suppose that the entrepreneur is Bayesian and knows the joint distribution of revenues and estimates. Thus, she would implement projects for which $E[\theta|s] \geq 0.4$. *Since* $E[\theta|s] = ps + 0.5(1-p)$, *if* $p > 0.2$, *the entrepreneur would set* *an implementation cutoff of* $0.5 - \frac{0.1}{p}$ *(i.e., she would implement only projects whose estimated returns are higher than the cutoff). If* $p \leq 0.2$*, then the entrepreneur would implement all projects.*

Now suppose that the entrepreneur takes the estimated returns at face value and, therefore, launches only projects whose estimated returns are higher than 0.4*. In the long run, if the entrepreneur examines the projects she implemented, she would observe that the estimated return is higher than the actual return on average. To see this, note that for any implementation cutoff* x*,*

$$E[s|s \geq x] = 0.5(x+1)$$

and

$$E[\theta|s \geq x] = 0.5p(x+1) + 0.5(1-p),$$

which implies

$$E[s-\theta|s \geq x] = 0.5x(1-p). \tag{1}$$

Plugging the cutoff $x = 0.4$ *into* (1) *yields an average difference of* $0.2(1-p)$ *between estimated and actual returns. Suppose that the entrepreneur resolves this discrepancy by adjusting her expectations downward by* $0.2(1-p)$*. Her expectation of* θ *would then be* $s - 0.2(1-p)$*. She would revise her strategy and set an implementation cutoff* $x = 0.4 + 0.2(1-p)$*. In the long run, increasing the implementation cutoff would increase the discrepancy in the entrepreneur's records (equation* (1)*). In turn, the increase in the long-run discrepancy would lead the entrepreneur to revise her strategy once again and, as a result, increase her implementation cutoff even further, and so on. This procedure converges to a steady state in which the entrepreneur sets an implementation cutoff of*

$$x = 0.4 + 0.5x(1-p).$$

Thus, there exists a steady state in which the entrepreneur launches projects if and only if $s \geq 0.8/(1+p)$*, observes a discrepancy of* $E[s-\theta|s \geq 0.8/(1+p)] = 0.4(1-p)/(1+p)$*, and expects a return of* 0.4 *at the cutoff.*

The entrepreneur in Example 1 sets a more conservative implementation cutoff than a Bayesian entrepreneur as $\frac{0.8}{1+p} > 0.5 - \frac{0.1}{p}$ for $p < 1$. The reason for this conservatism is that the entrepreneur bases her decisions on information from instances in which she implemented the projects; in these instances the estimated returns are, on average, higher than the actual returns. In an attempt to account for this discrepancy, the entrepreneur sets an excessively high implementation cutoff. As equation (1) illustrates, setting a more conservative implementation cutoff results in a larger discrepancy in the dataset that will be available in the future, which in turn results in an even more conservative implementation cutoff, and so on.

In our model, a decision maker (DM) employs the procedure similar to the one used by the entrepreneur in a more general setting. The DM faces a sequence of similar decision problems. In each problem, the DM observes an independent signal $s \in \mathbb{R}$ about the state of a relevant economic variable $\theta \in \mathbb{R}$ and then takes an action. One interpretation of the signal is a noisy estimate of the state's value. We assume that the higher the value of the signal is, the higher the expected value of the state, and that the higher the expected value of the state is, the higher the DM's optimal action.

The DM in our model does not know the joint distribution of signals and states. Instead, she uses a dataset that consists of past signals and state realizations to calibrate a simple prediction rule. Specifically, according to the DM's model of the world, the signal s reflects the true state θ up to a fixed bias b, and so she considers only additive prediction rules of the form[1] $\hat{\theta}(s) = s - b$. The DM calculates b by minimizing the quadratic loss function $E[(\theta - \hat{\theta})^2 | \mathcal{D}]$, which implies that b is the average difference between signals and state realizations in the dataset, as in Example 1. When facing a new decision, the DM uses the current signal s and the prediction rule, and chooses an action as if the state is $\hat{\theta}(s)$.

The DM's data is endogenous. After taking an action, the DM observes the state's realization with a probability that is increasing in her action. This monotonicity assumption is naturally satisfied in the investment decision setting. Another example is a second-price IPV auction in which the DM relies on a signal about the object's value. The higher the signal she observes, the higher the expected value of the object and the optimal bid are, and therefore the more likely the DM is to win the object and learn its actual value. The DM records all the state's realizations she observes and their respective signals in the dataset.

Since the DM's actions affect the data she observes and vice versa, we take an equilibrium approach to characterizing the steady state of this system. An equilibrium

[1]In Section 6, we discuss alternative families of prediction rules and their implications for our results.

consists of a strategy (a mapping from signals to actions) and a bias b, such that the strategy is a best response given the prediction $\hat{\theta}$, and the bias b minimizes the quadratic loss function given the data generated by the DM's strategy. We show that the DM's behavior converges to an equilibrium under mild conditions. Moreover, in equilibrium, the bias is always positive and, as a result, leads to a conservative interpretation of the signals as illustrated in Example 1.

We refer to the mapping from the DM's actions to the probability that she observes the actual realizations ex post as a *feedback function*. In Example 1, the feedback function assigns a probability of 1 to observing the returns of implemented projects and a probability of 0 to observing unimplemented ones. More generally, this function reflects basic properties of the environment in which the DM operates; for instance, different feedback functions may represent different types of organizational memory. We derive a tight condition that enables us to rank different feedback functions in terms of the equilibrium bias that they induce. Essentially, a feedback function ϕ induces a higher bias than the feedback function ϕ' for every objective distribution of states and signals if and only if ϕ dominates ϕ' in the likelihood ratio sense. In Section 5 we apply this condition to the setting of a second-price IPV auction in which bidders rely on a signal to bid and observe the actual value of the object only if they win. We use the fact that the feedback function depends on the number of bidders to show that the feedback-ranking condition implies that the bidders' bias is increasing in the number of bidders for any information structure.

We then analyze the externalities that DMs impose on one another. In this analysis, we interpret the DM as a sequence of DMs facing similar problems and assume that a fraction of the DMs use our misspecified calibration heuristic while the others are Bayesian in the sense that they know the joint distribution of states and signals and apply the Bayesian methodology. Since they know the joint distribution, Bayesian agents ignore the dataset when taking an action. Nonetheless, their actions select state realizations into the dataset, thereby affecting the information naive DMs rely on. We find that the presence of Bayesian DMs mitigates the externalities naive DMs impose on their successors through the data. Interestingly, the naive DMs' bias does not vanish even when the share of Bayesian DMs approaches 1.

Overall, our results indicate that calibrating a prediction rule in an attempt to account for the discrepancy in the data can lead to a seemingly pessimistic behavior. This may explain well-documented phenomena such as the hurdle rate premium puzzle, which is the tendency of firms to set investment hurdle rates that are substantially higher than the cost of capital (Poterba and Summers, 1995; Meier and Tarhan, 2007). This microfoundation for irrational pessimism in equilibrium is a contribution to the literature on misspecified beliefs. While the literature provides much evidence and many models of overoptimism, the literature has devoted little attention to irrational pessimism, which is an equally important topic: just like optimism, pessimism may lead individuals to erroneous conclusions and suboptimal choices.

Related literature

The apparent pessimism in our model is related to Compte's (2002) and Compte and Postlewaite's (2019) *cautious behavior*. While the apparent pessimism in our model follows from constraints on the DM's prediction rules (i.e., on her beliefs), the cautious behavior in Compte and Postlewaite (2019) follows from constraints on the agents' strategy space (e.g., restricting bidders to using a fixed shading factor across different signals). Essentially, our DM has coarse beliefs whereas their agents are restricted to using coarse strategies. Moreover, under both modeling approaches, the agents optimize an objective function given a certain dataset, while biases in the dataset can have an effect on their decision making. However, a key difference between the two approaches in this respect is that our DM chooses the prediction rule that has the best fit with the data, whereas their agents choose the (coarse) strategy that maximizes their payoff. Importantly, the data in our model is endogenous and affected by the DM's actions, which plays a key role in magnifying the DM's biased predictions. By contrast, in Compte and Postlewaite (2019) the dataset is exogenous.

The effects we find are also reminiscent of choice-driven optimism (Van den Steen, 2004), and the optimizer's curse (Smith and Winkler, 2006). These authors consider the perspective of an outside observer. By contrast, we consider the perspective of a DM who tries to account for selection bias while her actions affect the selection of state realizations into the dataset, which requires an equilibrium approach that is absent from the aforementioned papers.

This paper contributes to the literature on naive learning from endogenously selected data. Esponda and Vespa (2018) and Barron et al. (2019) provide empirical evidence that decision makers tend to neglect selection and extrapolate naively from endogenous data. Esponda (2008) shows that selection neglect can exacerbate adverse-selection problems and Jehiel (2018) lays the equilibrium microfoundations of overoptimism when there is positive selection in an investment decision setting similar to that of Example 1. In Section 5, we explain why Jehiel's procedure leads to opposite results from ours.

This paper belongs to a growing literature on decision-making and strategic interaction when agents hold misspecified models of the world (Piccione and Rubinstein, 2003; Eyster and Rabin, 2005; Jehiel, 2005; Esponda, 2008; Esponda and Pouzo, 2016; Spiegler, 2016; Heidhues et al., 2018, 2023).[2] For excellent reviews on this topic see Jehiel (2020) and Spiegler (2020). The calibration approach in this paper is related to Esponda and Pouzo's (2016) approach to learning with misspecified models. They propose a solution concept, called the Berk–Nash equilibrium, in which the DM's beliefs minimize the relative entropy with respect to her (misspecified) model of the world. Their model captures Bayesian learning under a misspecified model of the world. By contrast, our solution concept requires that observed discrepancies between signals and state realizations be resolved by fitting the mean of the signals to the mean of the states.

The paper proceeds as follows. We present the model in Section 2 and analyze it in Section 3. In Section 4 we show that the DM's behavior converges to an equilibrium. In Section 5 we use our results to study second-price IPV auctions and investment decision problems. Section 7 discusses alternative families of prediction rules, and Section 7 concludes. All proofs are relegated to the appendix.

2 The Model

The environment is composed of two random variables, namely, the state of nature θ and a signal s, which are distributed according to a bivariate probability distribution function $F(\theta, s)$ defined over $\mathbb{R}^2$. We assume that the (marginal) distribution over signals has a log-concave density and denote it by $f(s)$. The signal is unbiased ex ante in the sense that[3] $E(\theta) = E(s)$. We assume that $E[\theta|s]$ and $s - E[\theta|s]$ are nondecreasing in s. These assumptions imply that there exists $s^\star$ such that $s^\star = E[\theta|s^\star]$. When the support of the distribution of signals is unbounded, we require that $lim_{s' \to \infty} E[\theta|s \geq s']$ be unbounded as well. These assumptions are satisfied by many prevalent information

[2] There are ample evidence that economic agents depart from the Bayesian methodology (for a comprehensive review, see Benjamin (2019).

[3] With minor adjustments our results and intuitions continue to hold when the signals are biased, i.e., $E(s) \neq E(\theta)$. Specifically, the equilibrium bias would consist of the bias in the data $E(s) - E(\theta)$ and an additional bias. The assumption that $E(s) = E(\theta)$ allows us to focus on the latter bias.

structures. In particular, the signal can be the sum of the state and a zero-mean noise term.[4]

A strategy is a mapping $a : \mathbb{R} \to A$ from signals to actions, where $A = [\underline{a}, \overline{a}] \subset \mathbb{R}$ is the set of actions. The DM's payoff $\pi(\theta, a)$ is a function of the state and her action. The action that maximizes the payoff given a state θ is denoted by $a^\star(\theta)$ and assumed to be unique up to a measure zero set of states. We assume that $a^\star(\theta)$ is weakly increasing in θ and nondegenerate.

The DM has access to an infinitely large dataset $\mathcal{D}$ that consists of signals and state realizations from past decisions. When facing a new decision problem, she uses this dataset to generate a prediction rule $\hat{\theta}(s)$, which she then uses to interpret the signal at her disposal. She restricts attention to a family of simple prediction rules from which she chooses the one that fits the data best. Specifically, the DM considers rules of the form $\hat{\theta}(s) = s - b$, where b is a scalar, and chooses the rule that minimizes the quadratic loss function over the dataset

$$E[(\theta - \hat{\theta}(s))^2 | \mathcal{D}]. \tag{2}$$

Minimizing (2) yields

$$b_{\mathcal{D}} = E[s - \theta | \mathcal{D}]. \tag{3}$$

The dataset $\mathcal{D}$ is formed endogenously. The probability that each pair (s, θ) is observed and recorded in $\mathcal{D}$ depends on the action the DM chooses given s. Formally, a pair (s, θ) is recorded with probability $\phi(a(s))$, where $\phi : A \to [0, 1]$ is a feedback function and $a(s)$ is the action the DM chooses given s. We assume that $\phi(\cdot)$ is nondegenerate and nondecreasing, and that $\phi(a^\star(s)) > 0$ for some signal s. If the DM's strategy $a(\cdot)$ is constant over time and the dataset is not empty, then the proportion of each signal value s in $\mathcal{D}$ in the long run is

$$\frac{f(s)\phi(a(s))}{\int_{-\infty}^{\infty} f(s)\phi(a(s))ds}. \tag{4}$$

[4]Formally, suppose that $s = \theta + \epsilon$, where ϵ and θ are independently drawn from log-concave density functions and that $E(\epsilon) = 0$. As per Efron (1965), $s - E[\theta | \theta + \epsilon = s]$ and $E[\theta | s]$ are nondecreasing in s.

Thus, we can write (3) as

$$b_{\mathcal{D}} = \frac{\int_{-\infty}^{\infty} f(s)\phi(a(s))(s - E[\theta|s])ds}{\int_{-\infty}^{\infty} f(s)\phi(a(s))ds}. \tag{5}$$

We refer to the bias that solves this problem as the DM's *perceived bias*. In order to treat also cases in which the dataset is empty (e.g., when $a(s) = 0$ for every s), we rewrite (5) as

$$\int_{-\infty}^{\infty} f(s)\phi(a(s))(s - E[\theta|s] - b_{\mathcal{D}})ds = 0. \tag{6}$$

Note that when $a(s) = 0$ for all s, any perceived bias is consistent with the dataset.

We now present our equilibrium notion. The first requirement in Definition 1 is that the dataset be generated by the DM's strategy, and that the prediction rule minimize the quadratic loss function (2). The second requirement is that the DM's strategy be optimal given her predictions.

Definition 1 *A strategy $a(\cdot)$ and a bias b constitute an equilibrium if the following conditions are met:*

1. *b and $a(\cdot)$ satisfy equation* (6)*.*

2. *$a(s) = a^{\star}(s - b)$ for every signal s.*

Note that the definition does not preclude the possibility that $\mathcal{D} = \emptyset$ in equilibrium. This may occur when $\phi(\underline{a}) = 0$ and there exists a large enough bias b such that $a^{\star}(s - b) = \underline{a}$ for every signal s in the support of f. In such a case, every $b' \geq b$ is consistent with (6) and is part of an equilibrium. We refer to equilibria in which $\mathcal{D} = \emptyset$ as corner equilibria and to equilibria in which $\mathcal{D} \neq \emptyset$ as internal equilibria.

Discussion: Modeling Assumptions

Model misspecification. Despite having access to a large database containing information about the distribution of signals and states, our agent makes incorrect predictions and hence potentially suboptimal choices. However, in equilibrium, her misspecified model of the world is consistent with the evidence she collects. Thus, in addition to the simplicity of her calibration heuristic (compared, for example, to Bayesian updating), a possible justification for using it is that she has no apparent reason to doubt it. The idea that agents maintain confidence in their worldview as long as it is not contradicted by the data they gather has its roots in models of conjectural equilibria (Battigalli and Guaitoli, 1997) and self-confirming equilibria (Fudenberg and Levine, 1993) and is by now common in the bounded rationality literature (e.g., Spiegler, 2016).

Coarse reasoning. The DM's reasoning is coarse in the sense that she debiases all signals by the same additive term. A DM whose reasoning is fine would debias every signal s by $s - E[\theta|s]$, and so her predictions would be correct. In the spirit of the partially cursed equilibrium (Eyster and Rabin, 2005), we can define a "partially coarse" prediction rule by $\hat{\theta}_{\psi}(s) = (1-\psi)E[\theta|s] + \psi(s - E(s-\theta|\mathcal{D}))$. Note that for $\psi = 0$ this prediction rule corresponds to the Bayesian forecast whereas for $\psi = 1$ it corresponds to our baseline model. All of the propositions in the paper remain valid for any $0 < \psi \leq 1$ and the proofs are essentially identical. One interpretation of this partial misspecification is of a DM who takes into account the "bias per signal" $s - E[\theta|s]$ and the overall bias in her dataset, and assigns some weight to each of them. In the appendix we show that the equilibrium bias is increasing in ψ. That is, the more weight the DM assigns to the overall bias in the data, the larger the equilibrium bias and the departure from the rational model.

Availability of past signals. A key modeling assumption is that past signals (or at least their average value) are recorded in a dataset to which the DM has access. This is a plausible assumption in situations in which the same DM faces a sequence of similar decision problems or when multiple DMs who belong to the same organization face a sequence of similar decision problems and share their data. This assumption is perhaps less plausible when there is no relation between the DMs. Jehiel (2018) studies investment problems in which past signals are unavailable to the entrepreneur and obtains opposite results from ours. We discuss this difference in detail in Section 5.

3 Analysis

In this section, we show that the equilibrium bias is always positive and compare the behavior of a naive DM to the behavior of a Bayesian DM. Next, we study how the naive DM's behavior changes with the model's primitives. At the end of this section, we investigate the externalities that naive DMs impose on their successors. We start by showing that an equilibrium exists.

Proposition 1 *An equilibrium exists.*

To gain intuition for this result, note that a perceived bias b pins down the DM's strategy $a^\star(s-b)$, which pins down a dataset $\mathcal{D}_b$. If $\mathcal{D}_b \neq \emptyset$, then the dataset pins down a unique perceived bias of $E[s-\theta|\mathcal{D}_b]$. If $\mathcal{D}_b = \emptyset$, then the data is consistent with any "perception" of the bias and, in particular, with b. In the latter case, b is clearly an equilibrium. The proof shows that whenever there is no such bias, then the former case has a fixed point, i.e., there exists a bias $b^\star$ such that

$$b^\star = E[s-\theta|\mathcal{D}_{b^\star}].$$

Note that the equilibrium bias in our model need not be unique. If there exists more than one internal equilibrium, then we often focus on the two internal equilibria with the minimal and maximal biases. Denote the minimal and maximal (internal) equilibrium biases by $\overline{b}$ and $\underline{b}$, respectively.

Having established that an equilibrium exists, we now turn to study its properties. We start with the direction of the perceived bias.

Proposition 2 *In every equilibrium the perceived bias is positive, i.e.,* $b \geq 0$.

To understand this result, recall that if the DM were to observe all signals and their respective state realizations, there would be no bias at all. However, the DM's dataset includes only a selected sample of such pairs. In particular, since $\phi(\cdot)$ and $a^\star(\cdot)$ are increasing, the dataset contains disproportionately more cases in which the signal is high. The assumption that $s - E(\theta|s)$ is nondecreasing in s implies that the dataset also contains disproportionately more cases in which the difference $s-\theta$ is high, which in turn leads to a perceived upward bias.

We now provide a definition that allows us to rank different feedback functions according to the degree of selection they induce in the data. We then use this definition to show how the degree of selection affects the perceived bias in equilibrium.

Definition 2 *The feedback function* ϕ *dominates the feedback function* $\tilde{\phi}$ *in the likelihood ratio sense if* $\frac{\phi(a)}{\tilde{\phi}(a)}$ *is nondecreasing in* a.

To illustrate this definition, let $A = [0,1]$ and note that the feedback function $\tilde{\phi}(a) = a^n$ dominates the feedback function $\phi(a) = a^m$ for $m < n$. In general, the dominating feedback function returns relatively more observations in which the signal is high and the DM takes high actions, and fewer observations in which the signal is low and the DM takes low actions, compared to the dominated feedback function.

The next result establishes that if a feedback function dominates another feedback function in the likelihood ratio sense, then in an internal equilibrium it induces a higher bias.

Proposition 3 *If the feedback function ϕ dominates the feedback function $\tilde{\phi}$ in the likelihood ratio sense, then $\underline{b}_\phi \geq \underline{b}_{\tilde{\phi}}$ and $\overline{b}_\phi \geq \overline{b}_{\tilde{\phi}}$.*

The dominant feedback function ϕ assigns relatively more weight to high actions. Thus, intuitively, it assigns more weight to high signals and, as a result, more weight to instances in which the signal is high relative to the actual state realization. The more weight a feedback function assigns to these instances relative to instances in which signals are low, the higher the DM's perceived bias.

Likelihood ratio dominance is a tight condition in the sense that, if ϕ does not dominate $\tilde{\phi}$, then there exist a distribution $F(\cdot,\cdot)$, a payoff function π, and an optimal strategy $a^\star$, such that $\tilde{\phi}$ induces a higher bias than ϕ in equilibrium. The intuition for this tightness is that if ϕ does not dominate $\tilde{\phi}$ in the likelihood ratio sense, then there is some interval $[a_l, a_h]$ on which $\phi|_{[a_l,a_h]}$ is dominated by $\tilde{\phi}|_{[a_l,a_h]}$ in the likelihood ratio sense. It is possible to find a distribution $F(\cdot,\cdot)$ that is concentrated on that interval, a payoff function π, and a strictly increasing function $a^\star$, such that the result of Proposition 3 is reversed. The following corollary summarizes this discussion.

Corollary 1 *If ϕ does not dominate $\tilde{\phi}$ in the likelihood ratio sense then there exist a distribution $F(\cdot,\cdot)$ and a function $a^\star$, such that $\underline{b}_\phi < \underline{b}_{\tilde{\phi}}$ and $\overline{b}_\phi < \overline{b}_{\tilde{\phi}}$.*

At this point, it is worth comparing the prediction of our naive DM, $\hat{\theta}(s) = s - b$, to the prediction of a Bayesian DM, $E[\theta|s]$. Our DM's prediction is higher if and only if $s - b > E[\theta|s]$, or $s - E[\theta|s] > b$. Recall that, by assumption, $s - E(\theta|s)$ is nondecreasing in s. In an internal equilibrium, b is a weighted average of $s - E[\theta|s]$ and, therefore, $b \in (inf(s - E(\theta|s)), sup(s - E(\theta|s)))$. Thus, in an internal equilibrium, there is a signal $\hat{s}$ such that our DM's prediction is higher than the Bayesian DM's prediction if and only if $s > \hat{s}$.

While our DM's predictions are higher than a Bayesian DM's predictions for high signals, they are, on average, lower than the Bayesian DM's predictions. This follows from the equilibrium bias being positive (Proposition 2). Our DM's average prediction is $\int_{-\infty}^{\infty} f(s)(s-b)ds = E(s) - b = E(\theta) - b$, whereas the Bayesian DM's average prediction is $\int_{-\infty}^{\infty} f(s)E[\theta|s]ds = E(\theta)$. Corollary 2 summarizes this discussion.

Corollary 2 *In an internal equilibrium, there exists a signal $\hat{s}$ such that the naive DM's prediction $s - b$ is higher than a Bayesian DM's prediction $E[\theta|s]$ if and only if $s > \hat{s}$. Moreover, on average, the naive DM's prediction is weakly lower than the Bayesian DM's prediction.*

So far, we have assumed that the dataset on which the DM bases her decisions is generated by her own actions or by the actions of other DMs who use the same heuristic. However, in some situations DMs are more sophisticated and know the joint distribution of states and signals. Such DMs can apply the Bayesian methodology to make use of the signals at their disposal. Since their actions select different signals into the dataset, they affect the discrepancy in the data and the perceived bias. In turn, this bias affects the naive DMs' behavior and the signals they select into the dataset.

We now incorporate this idea into our model by assuming that a share $\alpha > 0$ of the DMs are Bayesian. Since Bayesian DMs are unaffected by the dataset (as they know the joint distribution of signals and states), varying their share enables us to study the externalities they impose on naive DMs without worrying that their own behavior is being affected by the presence of the naive DMs. In particular, it allows us to explore the implications for the naive DMs' equilibrium bias and whether the presence of Bayesian DMs brings the naive DMs' behavior closer to Bayesian behavior.

Throughout the analysis, we assume that a Bayesian DM plays the strategy[5] $a^{\star}(E(\theta|s))$. By similar arguments to those used in the proofs of Propositions 1 and 2, an equilibrium exists and the perceived bias is nonnegative for any $\alpha \in [0,1]$. Denote the highest and lowest equilibrium biases by $\overline{b}_{\alpha}$ and $\underline{b}_{\alpha}$, respectively. Next, we show that the presence of Bayesian DMs mitigates the discrepancy in the data and results in a lower bias: the more Bayesian DMs there are, the lower the bias is in an internal equilibrium.

Proposition 4 *In an internal equilibrium, both $\overline{b}_{\alpha}$ and $\underline{b}_{\alpha}$ are weakly decreasing in α.*

Relative to our naive DMs, Bayesian DMs assign a lower weight to the value of the signal as they take their prior beliefs into account. Therefore, as established in

[5]This assumption holds in various settings and, in particular, in both of our applications, namely, project selection and second-price IPV auctions.

Corollary 2, Bayesian DMs play lower (resp., higher) actions when the signal is high (resp., low). Since $\phi(\cdot)$ is increasing in a, it follows that actions taken by a Bayesian DM generate less (resp., more) feedback in situations in which the signal is high (resp., low). As a result, the average bias in the dataset is lower when the dataset contains more actions taken by Bayesian DMs. Thus, the presence of Bayesian DMs imposes a positive externality on naive DMs that leads them to choose lower actions. In turn, this equilibrium effect further decreases the bias in the data they rely on.

When the share of naive DMs vanishes, the bias in the data does not disappear. Indeed, $b_1 > 0$ except in the degenerate case where $\phi(a^\star(s))$ is constant for all s. This happens because Bayesian agents also take higher actions given higher signals, which implies that higher signals are more likely to be recorded in the dataset. Thus, decision-making based on data generated by Bayesian DMs results in a bias, albeit a smaller bias than the one obtained when DMs use the naive calibration procedure.

4 Convergence

Our analysis so far has focused on the steady state of the DM's learning process. We now describe the learning process in more detail and show that indeed it converges to a steady state. We assume that, in each period $t \in \mathbb{N}$, the DM receives a signal s_t and forms a prediction $\hat{\theta}_{\mathcal{D}_t}(s_t)$ based on the signal and the data obtained up to period t, $\mathcal{D}_t$. She then chooses the optimal action given the predicted state, $a^\star(\hat{\theta}_{\mathcal{D}_t}(s_t))$. At the end of the period, with probability $\phi(a^\star(\hat{\theta}_{\mathcal{D}_t}(s_t)))$, the signal s_t and the realization θ_t are recorded in the dataset that will be available in the future. For every period $t \in \mathbb{N}$, let $b_t = E[s - \theta | \mathcal{D}_t]$ if $\mathcal{D}_t \neq \emptyset$ and let $b_t = 0$ otherwise. As in the baseline model, $\hat{\theta}_{\mathcal{D}_t}(s_t) = s_t - b_t$. For simplicity, we assume that the support of $F(\cdot, \cdot)$ is bounded.

We now show that the sequence of biases converges in probability. Clearly, if it converges, it must converge to an equilibrium. That is, either b_t converges to a bias b for which $T(b) = b$ (an internal equilibrium) or to a bias b such that $\phi(a^\star(s - b)) = 0$ for every signal s (a corner equilibrium).

Proposition 5 *The sequence of biases* $(b_t)_{t=1}^{\infty}$ *converges in probability.*

To gain intuition for this result, note that, given a bias b_t, the expected difference between the next state and the signal that will be recorded in the DM's dataset is $T(b_t)$ in the next period. Roughly speaking, the proof shows that, despite the noisy process, for a sufficiently large t, if $T(b_t) > b_t$ then the bias must increase over time. Symmetrically, for a sufficiently large t, if $T(b_t) < b_t$, then b_t must decrease over time. Thus, it must converge to a bias b such that $T(b) = b$ or to a corner solution.

The proposition is silent as to which of the equilibria the sequence converges to. As an illustration, we simulated the setting of Example 1 with $p = 0.25, 0.5, 0.75$, and $c = 0.4$. For each value of p, we repeated the exercise $10,000$ times. The DM's behavior converged to a corner equilibrium (i.e., aborting all projects) in 1433 cases, 840 cases ,and 356 cases, respectively.

5 Applications

We now apply our results to two settings: investment decisions and second-price IPV auctions. In both applications, our naive calibration procedure leads to conservative behavior relative to the behavior of a Bayesian DM: rejection of marginally good projects in investment decisions and underbidding in auctions.

5.1 Investment Decisions

An entrepreneur selects which projects to implement based on their predicted revenue. Denote the revenue by θ, and the signal that reflects an estimate of this revenue by s. Implementing a project entails a cost of c. Denote a decision to implement a project by $a = 1$ and a decision to forgo it by $a = 0$. The entrepreneur wishes to implement a project if and only if $\theta \geq c$, and so $a^\star(\theta) = 1$ if $\theta \geq c$ and $a^\star(\theta) = 0$ otherwise.

The entrepreneur bases her decisions on a dataset that includes signals and actual revenues of implemented projects, i.e., $\phi(a) = a$. These projects were implemented by a set of entrepreneurs of which a share α are Bayesian and a share $1 - \alpha$ use our heuristic, where $0 \leq \alpha < 1$. Denote the entrepreneur's bias in an internal equilibrium by b_α. The naive entrepreneur predicts a revenue of $s - b_\alpha$ and, therefore, implements projects whose signals are higher than $c + b_\alpha$. Hence, in equilibrium, she uses an implementation cutoff of

$$s_\alpha = c + b_\alpha. \tag{7}$$

As a benchmark, note that a Bayesian entrepreneur implements a project if and only if $E[\theta|s] \geq c$. Since $E[\theta|s]$ is nondecreasing in s, there is a cutoff s_B such that a Bayesian entrepreneur implements a project if and only if $s \geq s_B$, where $E[\theta|s = s_B] = c$.

The next result characterizes the internal equilibrium bias as a function of[6] α.

Claim 1 *There exists at most one internal equilibrium. In such an equilibrium, (i) $s_\alpha \geq c$, (ii) $s_\alpha \geq s_B$, and (iii) s_α is decreasing in α.*

Parts (i) and (iii) follow from Propositions 2 and 4, respectively. Part (ii) shows that the naive entrepreneur's cutoff is higher than the Bayesian entrepreneur's cutoff although the naive entrepreneur's prediction is more optimistic than the Bayesian entrepreneur's prediction for sufficiently high signals (Corollary 2). To obtain intuition for this effect, note that since $E[\theta|s = s_B] = c$, the Bayesian entrepreneur's correction at the cutoff is $E[s-\theta|s = s_B]$. Were the naive entrepreneur to use the Bayesian cutoff s_B, then her bias would be $E[s - \theta|s > s_B]$, which is larger than $E[s - \theta|s = s_B]$ as $E[s - \theta|s]$ is nondecreasing. Thus, the naive entrepreneur would like to set a higher cutoff.

Note that the restriction to additive prediction rules of the form $\hat{\theta} = s - b$ does not impose a constraint on the entrepreneur's ability to maximize her payoff. To see this, note that the naive entrepreneur's strategy is identical to a Bayesian entrepreneur's strategy if she uses the prediction rule $\hat{\theta}(s) = s - E[s - \theta|s = s_B]$. In this case, both of them launch projects if and only if the signal is higher than s_B. The reason that restricting attention to additive prediction rules entails no loss of generality is the combination of the binary action space and the monotonicity of $a^\star(\cdot)$ and $E[\theta|s]$. We conclude that it is not the restriction to additive prediction rules that limits the DM, but rather it is her choice of rule that minimizes the quadratic loss and the dataset's endogeneity.

The result of Claim 1 makes clear that the bias of our DM in equilibrium, $E[s-\theta|s \geq s_\alpha]$, is higher than the optimal bias $E[s-\theta|s_B]$. This follows from the fact that $s_B \leq s_\alpha$ (Claim 1) and the assumption that $s - E[\theta|s]$ is nondecreasing. A related question is whether the DM's bias is payoff-enhancing relative to no bias at all. The answer depends on the parametric assumptions. For instance, suppose that $c = 1$, $\alpha = 0$, and $s = \theta + \epsilon$, where θ and ϵ are drawn independently from the standard normal distribution. A DM who takes signals at face value would set a cutoff of 1 while our DM would set a cutoff slightly lower than 2.34 in equilibrium. On average, projects whose estimated returns are between 1 and 2.34 yield less than 1, and so the DM's

[6] Clearly, in every corner equilibrium, all projects are forgone.

bias is payoff-enhancing in this case.[7] If we change the cost to $c = 0$, then our DM would set a cutoff higher than 0 in equilibrium, whereas a DM who takes signals at face value would set a cutoff of 0, similar to that of a Bayesian DM. Thus, the DM's bias is payoff-degrading in this case.

The results in this section run counter to the results obtained in Jehiel's (2018) model of investment decisions even though the setting is similar. In particular, in the equilibrium of his model, the entrepreneur uses an implementation cutoff that is too low relative to the cutoff of a Bayesian entrepreneur. It is instructive to discuss the different modeling assumptions that lead to this difference.

In both models, the entrepreneur has access to the realized revenues of all of the projects that were implemented in the past and a signal about the current project. While in our model the DM has access to the original signals received when the past projects were implemented, in Jehiel's model the entrepreneur obtains a *new* signal about each of the past projects. Jehiel's entrepreneur believes that the current project's revenue will be identical to the average revenue of past projects for which he received the signal he received for the current project. Jehiel shows that if the entrepreneur's signals are independent of one another, then, in equilibrium, the entrepreneur sets an implementation cutoff that is too low. The intuition for this overoptimism is that the dataset that is used to evaluate the current project includes only implemented projects, i.e., projects for which previous entrepreneurs received a signal above their implementation cutoff. The expected revenue of these projects is higher than the expected revenue of a new project conditional on receiving the same signal. However, the entrepreneur ignores this selection and, to use our terminology, believes that her signal is downward biased and adjusts her implementation cutoff downwards.

In both Jehiel's model and ours, the data that the entrepreneur uses to evaluate new projects suffers from selection bias. However, while Jehiel's entrepreneur uses a dataset about projects whose realized revenue is relatively high conditional on the signals, our entrepreneur uses a dataset about projects whose signals are relatively high conditional on the realized revenue. Thus, while Jehiel's entrepreneur concludes that the signals are downward biased, our entrepreneur concludes that the signals are upward biased. It follows that the overoptimism found in Jehiel (2018) would remain if the DM used the dataset to form a prediction about the current project as our DM (i.e., if she were to calculate the average bias of all past projects). Thus, the availability of past signals

[7] A Bayesian DM would set a cutoff of 2 under these parameters.

determines whether the DM's prediction is optimistic or pessimistic.

Other applications of our model: Medical referrals, recommendation systems, and credit provision

While we have been using the investment decision terminology, the analysis in this section is relevant in other contexts as well. For example, the DM might be a physician who refers patients for a diagnostic test based on the result of a screening test (e.g., an Antigen test followed by a PCR test for Covid 19).[8] To evaluate the screening test score, she uses data that includes results (of both tests) of past patients. Alternatively, the DM might be an individual who uses a recommendation system to decide which products and services to consume and naively calibrates the recommendation she receives based on her actual enjoyment in previous situations in which she consumed the product. Finally, the DM might be a credit officer who approves credit applications based on a credit score that is calibrated based on the return rate of previous successful applications but not of unsuccessful ones. In all of these situations our results imply a seemingly pessimistic behavior: setting the bar too high for medical referrals and credit provision, and downgrading positive recommendations from an algorithm or a friend.

5.2 Second-Price IPV Auctions

While our baseline model considers the decision of a single DM (or a sequence of such DMs), its framework can be extended to strategic situations. In general, this requires extending the payoff and feedback functions, and making assumptions on how agents reason about other agents' behavior. However, when a game is dominance-solvable, as in a second-price IPV auction, such assumptions become moot.

We assume that there are $n \geq 2$ bidders who bid for an object, each of whom receives a signal s_i about the value she will derive from the object, θ_i. The value and its signal are independently drawn from $F(\theta, s)$ for each bidder. Recall that, in a second-price IPV auction, bidding one's value is a dominant strategy, i.e., $a^\star(\theta_i) = \theta_i$. To bypass the problem of potentially negative bids, we assume that $supp(F) = [1, 2]^2$. Denote bidder i's perceived bias by b_i. Bidder i's predicted value is $\hat{\theta}(s_i) = s_i - b_i$.

[8] A screening test provides information about the risk of a certain disorder or condition. It is typically less costly and so it is used by a larger group of patients. A diagnostic test establishes the existence of a condition or a disease.

We assume that each bidder i uses the bidding function $a(s_i) = a^\star(\hat{\theta}(s_i)) = a^\star(s_i - b_i)$. Finally, a bidder learns her true valuation of the object if and only if she wins the object, i.e., $\phi(a_i, a_{-i}) = 1$ if $a_i > max_{j \neq i} a_i$ and $\phi(a_i, a_{-i}) = 0$ otherwise.

Following is the formal definition of an equilibrium in this game, which extends Definition 1.

Definition 3 *An equilibrium in a second-price IPV auction is a profile of bidding functions such that (i) the entire profile constitutes a Nash equilibrium in undominated strategies and (ii) each bidder's bidding function is part of an equilibrium at the individual level.*

In a symmetric (internal) equilibrium, a bidder obtains feedback (wins the object) if and only if her signal is the highest; i.e., she receives feedback with probability $F(s_i)^{n-1}$. This feedback function is independent of the bidders' bias, and so there is a unique bias that is consistent with our calibration procedure.

Next, we establish that the equilibrium is unique and that its bias is strictly positive. Moreover, it provides comparative statics with respect to the number of bidders.

Claim 2 *There exists a unique symmetric internal equilibrium. In such an equilibrium it holds that $a(s) < s$. Moreover, $a(s)$ is decreasing in n.*

In a symmetric equilibrium, the bias is strictly positive. This follows from the strict monotonicity of the bidding function. As a result, different signals lead to different frequencies of observing the actual realization, which precludes the possibility of a solution in which the bias is null. The comparative statics with respect to the number of bidders follow directly from Proposition 3. To see this, recall that the feedback probability when there are n bidders is $F(s)^{n-1}$. This function is dominated in the likelihood ratio sense by $F(s)^{m-1}$ for $m > n$ bidders.

As in a second-price IPV auction with Bayesian bidders, the bidder with the highest signal wins the object. Thus, the equilibrium outcome is efficient. However, the naive calibration procedure affects the bidding strategy and, therefore, the division of surplus.

We now turn to the auctioneer's perspective and compare her expected revenue in the case in which bidders are Bayesian to the case in which they use our heuristic.

Claim 3 *The auctioneer's revenue when bidders use the naive calibration procedure is lower than her revenue when bidders are Bayesian.*

This comparison is not obvious ex ante as our bidders' bids can be higher or lower than the ones submitted in a second-price auction with Bayesian bidders (Corollary 2). Denote the highest and second-highest signals by $s_{(n)}$ and $s_{(n-1)}$, respectively. When bidders are naive, the winner pays $s_{(n-1)} - b$, where the bias b is the average difference between the highest signal and the expected value conditional on receiving the highest signal. Thus, a naive winner pays, on average, $E(s_{(n-1)}) - (E(s_{(n)}) - E(\theta|s_{(n)}))$. When bidders are Bayesian, the winner pays the expected value of the object given the second-highest signal, $E(\theta|s_{(n-1)})$. Since $s - E(\theta|s)$ is increasing in s, and the distribution of $s_{(n)}$ first-order stochastically dominates the distribution of $s_{(n-1)}$, it holds that $E(s_{(n-1)} - E(\theta|s_{(n-1)})) < E(s_{(n)} - E(\theta|s_{(n)}))$. Therefore, on average, the naive winner bids less than the Bayesian one.

A corollary of Claim 3 is that, in expectation, if all bidders are naive, they obtain a higher payoff than they would obtain if they were all Bayesian. However, at an individual level (where the strategies of the other bidders are held fixed), as in the investment decision setting, the equilibrium prediction rule need not be optimal, even within the class of prediction rules of the form $\hat{\theta}(s) = s - b$. Unlike in the investment decision setting, in the auction setting the restriction to additive prediction rules entails a loss. The optimal bid given complete knowledge of the joint distribution of signals and states is $E[\theta|s]$. However, a fixed bias b that satisfies $s - b = E[\theta|s]$ does not exist.

Our analysis of second-price auctions is related to the winner's curse in IPV auctions identified by Compte (2002). In his model, bidders in a procurement auction rely on a noisy estimate of the cost. Due to selection bias, the estimate is likely to be lower than the actual cost conditional on winning. Compte analyzes the model under an assumption that bidders are constrained to using a coarse bidding function: they choose a fixed markup and add it to their estimated cost to correct for the selection bias. The bidders in Compte (2002) maximize their net payoff subject to the above constraint. In equilibrium, the markup is positive (to correct for selection) and increasing in the number of bidders. The intuition for the latter effect is that, as in our model, competition exacerbates the selection bias (i.e., the highest signal is expected to have a greater positive noise when there are more bidders).

It is possible to apply our model to procurement auctions and show that the bidders in our model adjust their bids more than the bidders in Compte (2002). The difference follows from the fact that bidders in both models have different objectives. While our bidders form a constrained optimal prediction rule given their model of the world, the bidders in Compte maximize their payoff subject to the fixed shading constraint. The overcorrection in our model is also related to a conceptual difference between the two models. In our model, bidders' behavior stems from a "personal" equilibrium that exacerbates the selection bias they are exposed to whereas in Compte (2002) bidders' shading is a standard (constrained) best response to the other bidders' behavior.

6 Alternative Prediction Rules

Throughout the analysis we examined the behavior of a DM who considers additive prediction rules of the form $\hat{\theta}(s) = s - b$. This family of additive prediction rules satisfies the following property: either $\hat{\theta}(s) \geq s$ for every s or $\hat{\theta}(s) \leq s$ for every s. In general, it is reasonable to focus on prediction rules that satisfy this property in situations where the DM is inclined to believe that all signals are biased in the same direction and is attempting to determine the direction and extent of such a bias. Our main insights apply to any family of rules that satisfy this condition. However, they may not hold when it is violated. To conclude the paper, we provide two examples of families of rules, one that meets the above property and one that does not.

The family of multiplicative prediction rules (i.e., rules of the form $\hat{\theta}(s) = \eta s$) satisfies the above property when all of the possible signal realizations are positive. These rules are simple to implement and are natural in various economic contexts. Let us consider a DM who chooses the rule that minimizes the quadratic loss function (2) for this family of rules. In a given dataset $\mathcal{D}$, this minimization yields $\eta_{\mathcal{D}} = E[\theta|\mathcal{D}]/E[s|\mathcal{D}]$ and the equilibrium condition becomes

$$\int_{-\infty}^{\infty} f(s)\phi(a^{\star}(\eta s))(\eta s - E[\theta|s])ds = 0.$$

The monotonicity of $a^{\star}(\cdot)$ and $\phi(\cdot)$ together with the monotonicity of $s - E[\theta|s]$ implies that $E[s|\mathcal{D}] \geq E[\theta|\mathcal{D}]$. Thus, $\eta_D \leq 1$, which is analogous to $b \geq 0$ in the additive model. Moreover, $\eta \leq 1$ implies that, on average, the DM's prediction is lower than the Bayesian prediction as in the additive model.[9] Thus, the main insights of the paper hold when the DM considers multiplicative prediction rules. In fact, all the results in Section 3 hold in this case.

[9]To see this, note that the average prediction of a DM who considers multiplicative prediction rules is $E[\hat{\theta}] = \int_{-\infty}^{\infty} f(s)s\eta_D ds = \eta_D E[s] \leq E[s] = E[\theta]$.

The family of linear prediction rules (i.e., rules of the form $\hat{\theta}(s) = \beta_0 + \beta_1 s$) does not necessarily satisfy the property discussed above,[10] which, for some parameters, yields opposite results from those of our baseline model. In the appendix, we illustrate this point using the investment decision setting. We show that if $E[\theta|s]$ is convex (resp., concave) then the DM chooses an excessively high (resp., low) implementation cutoff.

7 Concluding Remarks

We studied a model in which an agent has a misspecified model of the world that she calibrates based on a dataset that suffers from selection bias. The agent inadvertently contributes to this selection bias twice: first by taking actions that affect the data collection procedure and second by trying to correct for the apparent bias. We show that the naive calibration procedure can generate substantial biases and exacerbate the misspecification errors. Our findings indicate that these errors are consistently in the same direction and result in a seemingly pessimistic behavior.

The naive calibration procedure may result in a lower payoff relative to no calibration at all. For instance, in the investment decision setting of Example 1, a DM who takes the estimates at face value implements more projects than our naive calibrator (i.e., she sets a lower cutoff), and all these additional projects are profitable in expectation. If one interprets the attempt to calibrate the signal rather than taking it at face value as an indication of sophistication, then this example illustrates that a higher degree of sophistication may actually lead to a worse outcome. In recent years, it has been shown that in strategic interactions, players who are more sophisticated may obtain lower payoffs (e.g., Ettinger and Jehiel, 2010; Eyster and Piccione, 2013). We thus contribute to the literature by showing that a higher degree of sophistication may worsen outcomes in decision problems.

The naive calibration procedure can also improve the DM's welfare. Not only can it yield a higher payoff than no calibration at all (as shown in Section 5.1), but also, when signals are provided by a strategic agent, the procedure can yield a higher payoff relative to the payoff obtained by a Bayesian DM. For instance, consider a buyer who receives noisy information about the suitability or quality of various products advertised by a strategic seller. The latter might have an incentive to send biased signals and to add

[10] This family of rules includes our baseline model, in which $\beta_1 = 1$, and multiplicative prediction rules, in which $\beta_0 = 0$. Under these restrictions, the above property is satisfied.

noise to the transmitted information. However, a buyer who uses the naive calibration procedure not only corrects for the bias, but also interprets the noise as an additional systematic upward bias and corrects for that as well, which lowers her willingness to pay. Thus, a strategic seller who takes the DM's calibration procedure into account has an incentive to provide her with the most precise information possible.

Our analysis has implications for real-world procedures that correct for optimism bias, such as reference class forecasting. These procedures typically consider the cost overruns in similar past projects and raise new cost forecasts accordingly. However, these procedures ignore the selection bias problem that is at the heart of our analysis. As long as the cost of such projects affects the decision to carry them out, our results imply that these adjustments may yield excessively high cost forecasts. Thus, the attempt to fully correct for cost overruns based on past projects may lead decision makers who rely on such forecasts to make suboptimal choices such as forgoing socially beneficial projects.

References

AACE International (2012): *Total Cost Management Framework–An Integrated Approach to Portfolio, Program, and Project Management.*

Bagnoli, M. and T. Bergstrom (2005): "Log-concave probability and its applications," *Economic Theory*, 26, 445–469.

Barron, K., S. Huck, and P. Jehiel (2019): "Everyday econometricians: Selection neglect and overoptimism when learning from others," Working Paper.

Battigalli, P. and D. Guaitoli (1997): "Conjectural equilibria and rationalizability in a game with incomplete information," *Decisions, Games, and Markets,*, 97–124.

Benjamin, D. J. (2019): "Chapter 2 - Errors in probabilistic reasoning and judgment biases," in *Handbook of Behavioral Economics - Foundations and Applications 2*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson, North-Holland, vol. 2 of *Handbook of Behavioral Economics: Applications and Foundations 1*, 69–186.

Brashear, J., A. Becker, and D. Faulder (2001): "Where have all the profits gone?" *Journal of petroleum technology*, 53, 20–73.

Chen, M. and J. Dyer (2009): "Inevitable disappointment in projects selected based on forecasts," SPE Journal, vol. 14, 216–221.

Compte, O. (2002): "The winner's curse with independent private values," *Working Paper*.

Compte, O. and A. Postlewaite (2019): *Ignorance and Uncertainty*, Cambridge University Press.

Efron, B. (1965): "Increasing properties of Polya frequency function," *The Annals of Mathematical Statistics*, 36, 272–279.

Erdos, P. (1949): "On a theorem of Hsu and Robbins," *The Annals of Mathematical Statistics*, 20, 286–291.

Esponda, I. (2008): "Behavioral equilibrium in economies with adverse selection," *American Economic Review*, 98, 1269–1291.

Esponda, I. and D. Pouzo (2016): "Berk–Nash equilibrium: A framework for modeling agents with misspecified models," *Econometrica*, 84, 1093–1130.

Esponda, I. and E. Vespa (2018): "Endogenous sample selection: A laboratory study," *Quantitative Economics*, 9, 183–216.

Ettinger, D. and P. Jehiel (2010): "A theory of deception," *American Economic Journal: Microeconomics*, 2, 1–20.

Eyster, E. and M. Piccione (2013): "An approach to asset pricing under incomplete and diverse perceptions," *Econometrica*, 81, 1483–1506.

Eyster, E. and M. Rabin (2005): "Cursed equilibrium," *Econometrica*, 73, 1623–1672.

Flyvbjerg, B. (2008): "Curbing optimism bias and strategic misrepresentation in planning: Reference class forecasting in practice," *European planning studies*, 16, 3–21.

Flyvbjerg, B., C. Glenting, and A. Rønnest (2004): "Procedures for dealing with optimism bias in transport planning," *London: The British Department for Transport, Guidance Document*.

Fudenberg, D. and D. K. Levine (1993): "Self-confirming equilibrium," *Econometrica: Journal of the Econometric Society*, 523–545.

Heidhues, P., B. Kőoszegi, and P. Strack (2023): "Misinterpreting Yourself," *Available at SSRN 4325160.*

Heidhues, P., B. Kőszegi, and P. Strack (2018): "Unrealistic expectations and misguided learning," *Econometrica*, 86, 1159–1214.

HM Treasury (2003a): *The Green Book: appraisal and evaluation in central government: Treasury guidance*, HM Treasury.

——— (2003b): "Optimism Bias: Supplementary Green Book Guidance," .

Hsu, P.-L. and H. Robbins (1947): "Complete convergence and the law of large numbers," *Proceedings of the national academy of sciences*, 33, 25–31.

Jehiel, P. (2005): "Analogy-based expectation equilibrium," *Journal of Economic Theory*, 123, 81–104.

——— (2018): "Investment strategy and selection bias: An equilibrium perspective on overoptimism," *American Economic Review*, 108, 1582–1297.

——— (2020): "Analogy-based expectation equilibrium and related concepts: Theory, applications, and beyond," Working Paper.

Kahneman, D. and D. Lovallo (1993): "Timid choices and bold forecasts: A cognitive perspective on risk taking," *Management science*, 39, 17–31.

Meier, I. and V. Tarhan (2007): "Corporate investment decision practices and the hurdle rate premium puzzle," *Available at SSRN 960161.*

Nesvold, E. and R. B. Bratvold (2022): "Debiasing probabilistic oil production forecasts," *Energy*, 258, 124744.

Piccione, M. and A. Rubinstein (2003): "Modeling the economic interaction of agents with diverse abilities to recognize equilibrium patterns," *Journal of the European Economic Association*, 1, 212–223.

Poterba, J. M. and L. H. Summers (1995): "A CEO survey of US companies' time horizons and hurdle rates," *MIT Sloan Management Review*, 37, 43.

Smith, J. E. and R. L. Winkler (2006): "The optimizer's curse: Skepticism and postdecision surprise in decision analysis," *Management Science*, 52, 311–322.

Spiegler, R. (2016): "Bayesian networks and boundedly rational expectations," *The Quarterly Journal of Economics*, 131, 1243–1290.

——— (2020): "Behavioral implications of causal misperceptions," *Annual Review of Economics*, 12, 81–106.

UK Department for Transport and Oxford Global Projects (2020): *Updating the Evidence Behind the Optimism Bias Uplifts for Transport Appraisals.*

Van den Steen, E. (2004): "Rational overoptimism (and other biases)," *American Economic Review*, 94, 1141–1151.

Appendix: Proofs

Proof of Proposition 1. Suppose that the DM plays the strategy $a^\star(s-b)$. If there exists a bias b such that $\phi(a^\star(s-b))=0$ for every signal s, then there exists an equilibrium in which the bias is b and $D=\emptyset$. Otherwise, the DM's dataset is nonempty and the DM's strategy induces a perceived bias of

$$T(b)=\frac{\int_{-\infty}^{\infty} f(s)\phi(a^\star(s-b))[s-E[\theta|s]]ds}{\int_{-\infty}^{\infty} f(s)\phi(a^\star(s-b))ds}. \tag{8}$$

To establish the existence of an equilibrium, we show that there exists a bias b for which $T(b)=b$. Note that $lim_{b\to-\infty}\phi(a^\star(s-b))=\phi(\overline{a})\in(0,1]$ for every s, and so

$$lim_{b\to-\infty}T(b)-b=\frac{\int_{-\infty}^{\infty} f(s)[s-E[\theta|s]]ds}{\int_{-\infty}^{\infty} f(s)ds}-b=E[s]-E[\theta]-b=\infty.$$

Since $s-E[\theta|s]$ is nondecreasing, $T(b)$ is no higher than

$$T^x(b):=\frac{\int_{x+b}^{\infty} f(s)\phi(a^\star(s-b))[s-E[\theta|s]]ds}{\int_{x+b}^{\infty} f(s)\phi(a^\star(s-b))ds}. \tag{9}$$

Moreover, $T^x(b)$ is nondecreasing in x. Fix a small $\epsilon>0$. There exists a sufficiently large x such that for every pair $s,s'>x$ it holds that $|\phi(a^\star(s))-\phi(a^\star(s'))|<\epsilon$. Hence, for such x it holds that $T^x(b)$ is arbitrarily close to $E[s-\theta|s \geq x+b]$ for every b.

Log-concavity of the signals' distribution implies that $E[s|s \geq x+b] - x - b$ is decreasing in b (Bagnoli and Bergstrom, 2005). Hence, for a sufficiently large b we have that $E[s|s \geq x+b] - b - E[\theta|s \geq x+b] < 0$. Thus, $E[s-\theta|s \geq x+b] < b$ and, therefore, $T(b) < T^x(b) < b$. By the intermediate value theorem there exists b for which $T(b) = b$.

Proof of Proposition 2. First, consider a corner equilibrium and recall that, by assumption, there exists a signal s such that $\phi(a^\star(s)) > 0$. Since $\phi(\cdot)$ and $a^\star(\cdot)$ are increasing, $\phi(a^\star(s-b)) > 0$ for every $b < 0$. But this is in contradiction to the equilibrium being a corner one.

Second, consider an internal equilibrium and recall that the equilibrium bias b satisfies $b = T(b)$ (see (5)). Let $s^\star$ denote an arbitrary signal that satisfies $s^\star - E[\theta|s^\star] = 0$. It follows that

$$\int_{-\infty}^{s^\star} f(s)\phi(a^\star(s-b))[s-E(\theta|s)]ds \geq \int_{-\infty}^{s^\star} f(s)\phi(a^\star(s^\star-b))[s-E(\theta|s)]ds$$

and

$$\int_{s^\star}^{\infty} f(s)\phi(a^\star(s-b))[s-E(\theta|s)]ds \geq \int_{s^\star}^{\infty} f(s)\phi(a^\star(s^\star-b))[s-E(\theta|s)]ds.$$

The sum of the RHS of the two inequalities is 0 as $E(\theta) = E(s)$. Thus, the numerator of $T(b)$, which is equal the sum of the LHS of the two inequalities, is nonnegative. Furthermore, its denominator is positive. Hence, $T(b) = b \geq 0$ in equilibrium.

Proof of Proposition 3. To prove this result, we show that $T_\phi(b) \geq T_{\tilde{\phi}}(b)$ for every b, where T_ϕ (resp., $T_{\tilde{\phi}}$) denotes the operator T when the feedback function is ϕ (resp., $\tilde{\phi}$). Observe that $T_\phi(b) \geq T_{\tilde{\phi}}(b)$ if and only if

$$(10)\quad \frac{\int_{-\infty}^{\infty} f(s)\phi(a^\star(s-b))[s-E(\theta|s)]ds}{\int_{-\infty}^{\infty} f(s)\phi(a^\star(s-b))ds} \geq \frac{\int_{-\infty}^{\infty} f(s)\tilde{\phi}(a^\star(s-b))[s-E(\theta|s)]ds}{\int_{-\infty}^{\infty} f(s)\tilde{\phi}(a^\star(s-b))ds}.$$

Without loss of generality, we can assume that the feedback functions satisfy

$$(11)\quad \int_{-\infty}^{\infty} f(s)\phi(a^\star(s-b))ds = \int_{-\infty}^{\infty} f(s)\tilde{\phi}(a^\star(s-b))ds.$$

Since $\frac{\phi(a^\star(s-b))}{\tilde{\phi}(a^\star(s-b))}$ is increasing in s, (11) implies that there exists $s^\star$ such that $\phi(a^\star(s^\star - b)) \geq \tilde{\phi}(a^\star(s^\star - b))$ for $s > s^\star$ and the inequality is reversed for $s < s^\star$. Using the normalization in (11), we can write (10) as

$$\int_{-\infty}^{s^\star} f(s)\left(\phi(a^\star(s-b)) - \tilde{\phi}(a^\star(s-b))\right)[s - E(\theta|s)]ds + \int_{s^\star}^{\infty} f(s)\left(\phi(a^\star(s-b)) - \tilde{\phi}(a^\star(s-b))\right)[s - E(\theta|s)]ds \geq 0. \tag{12}$$

Since $s - E(\theta|s)$ is nondecreasing in s, the LHS of (12) is higher than

$$\int_{-\infty}^{s^\star} f(s)\left(\phi(a^\star(s-b)) - \tilde{\phi}(a^\star(s-b))\right)[s^\star - E(\theta|s^\star)]ds + \int_{s^\star}^{\infty} f(s)\left(\phi(a^\star(s-b)) - \tilde{\phi}(a^\star(s-b))\right)[s^\star - E(\theta|s^\star)]ds, \tag{13}$$

which, by (11), is equal to zero.

Proof of Proposition 4. In an internal equilibrium, $T_\alpha(b) = b$, where

$$T_\alpha(b) := \frac{\int_{-\infty}^{\infty} f(s)[(1-\alpha)\phi(a^\star(s-b)) + \alpha\phi(a^\star(E[\theta|s])](s - E[\theta|s])ds}{\int_{-\infty}^{\infty} f(s)[(1-\alpha)\phi(a^\star(s-b)) + \alpha\phi(a^\star(E[\theta|s])]ds}.$$

Note that $T_\alpha(b) = b$ if and only if

$$\int_{-\infty}^{\infty} f(s)g_\alpha(s)(s - b - E[\theta|s])ds = 0, \tag{14}$$

where

$$g_\alpha(s) = (1-\alpha)\phi(a^\star(s-b)) + \alpha\phi(a^\star(E[\theta|s])).$$

Assume that $T_\alpha(b^\star) = b^\star$, i.e., $b^\star$ is part of an equilibrium. Consider a signal $s^\star$ such that $s^\star - b^\star = E[\theta|s^\star]$. Note that $g_\alpha(s^\star)$ is independent of α. Moreover, By the monotonicity of $\phi(a(\cdot))$, it follows that $g_\alpha(s)$ is increasing (resp., decreasing) in α for every $s < s^\star$ (resp., $s > s^\star$). Since $s - b^\star - E[\theta|s] \leq 0$ if, and only if, $s \leq s^\star$, the LHS of (14) is nonpositive for $\alpha' > \alpha$. That is, $T_{\alpha'}(b^\star) \leq b^\star$ for $\alpha' > \alpha$. Similarly, $T_{\alpha'}(b^\star) \geq b^\star$ for $\alpha' < \alpha$.

Fix α and consider $\underline{b}_\alpha$. Let $\alpha' > \alpha$. Since $T_{\alpha'}(\underline{b}_\alpha) \leq \underline{b}_\alpha$, $T_{\alpha'}(0) \geq 0$, and $T_{\alpha'}(\underline{b}_{\alpha'}) = \underline{b}_{\alpha'}$ by continuity, it holds that $0 \leq \underline{b}_{\alpha'} \leq \underline{b}_\alpha$. Next, consider $\bar{b}_\alpha$. By previous arguments,

$T_\alpha(\bar{b}_{\alpha'}) \geq \bar{b}_{\alpha'}$. Note that $T_\alpha(b) < b$ for $b > \bar{b}_\alpha$ since $lim_{b\to\infty} T_\alpha(b)$ is finite. Therefore, $\bar{b}_{\alpha'} \leq \bar{b}_\alpha$.

Proof of Claim 1. Note that when $\alpha = 1$ it holds that $\mathcal{D} = \{(s,\theta)|s \geq s_B\}$. Moreover, there exists a unique b_1 that solves $b_1 = E[s-\theta|s \geq s_B]$. Since $E(s) = E(\theta)$ and $E[s-\theta|s]$ is nondecreasing, it follows that $b_1 \geq 0$. Since $c = s_B - E[s-\theta|s = s_B]$, it follows that

$$s_B = c + E[s-\theta|s = s_B] \leq c + E[s-\theta|s \geq s_B] = c + b_1 = s_1.$$

By Proposition 4, $\underline{b}_1 \leq \underline{b}_\alpha$ for any $\alpha < 1$. Hence, $c + \underline{b}_\alpha \geq c$ and $c + \underline{b}_\alpha \geq s_B$ for any α.

We now assume that b_α is part of an internal equilibrium and show that there cannot exist another internal equilibrium with a perceived bias $b'_\alpha > b_\alpha$. Let γ be the share of observations in the data there were induced by Bayesian DMs in the equilibrium in which the bias is b_α. That is,

$$\gamma_{b_\alpha} = \frac{\alpha(1-F(s_B))}{\alpha(1-F(s_B)) + (1-\alpha)(1-F(c+b_\alpha))}.$$

Note that $\gamma_{b'_\alpha} > \gamma_{b_\alpha}$ for $b'_\alpha > b_\alpha$. In an internal equilibrium,

$$T(b_\alpha) - b_\alpha = \gamma_{b_\alpha}(E[s-\theta|s \geq s_B] - b_\alpha)) + (1-\gamma_{b_\alpha})(E[s-\theta|s \geq c + b_\alpha] - b_\alpha) = 0.$$

Moreover, $E[s-\theta|s \geq c+b_\alpha] \geq E[s-\theta|s \geq s_B]$, and so

$$\gamma_{b'_\alpha}(E[s-\theta|s \geq s_B] - b_\alpha)) + (1-\gamma_{b'_\alpha})(E[s-\theta|s \geq c + b_\alpha] - b_\alpha) \leq 0.$$

Due to the log-concavity of the signals' distribution, $E[s-\theta|s \geq c+b_\alpha] - b_\alpha$ is decreasing in b_α. Clearly, $E[s-\theta|s \geq s_B] - b_\alpha$ is strictly decreasing in b_α. We conclude that $T(b'_\alpha) - b'_\alpha < 0$, in contradiction to b'_α being part of an equilibrium.

Proof of Claim 2. Since agents' strategies are symmetric, we can write the bias as

(15) $$b = \frac{\int_1^2 f(s)F(s)^{n-1}[s - E(\theta|s)]ds}{\int_1^2 f(s)F(s)^{n-1}ds}.$$

As the RHS of (15) is independent of b, it has a unique solution. Hence, these exists a unique internal equilibrium. Since (i) $E(s) = E(\theta)$, (ii) $s - E(\theta|s)$ is nondecreasing and nondegenerate, and (iii) $F(s)$ is *strictly* increasing, it follows that $b > 0$ and, therefore, $a(s) = a^{\star}(s - b) = s - b < s$. Since $\phi_n(a^{\star}(s - b_n)) = F(s)^{n-1}$, it follows that $\frac{\phi_{n+1}(\cdot)}{\phi_n(\cdot)} = F(s)$ is increasing in s and, by Proposition 3, the proof is complete.

Proof of Claim 3. Let $s_1, s_2, ..., s_n$ be a random sample of signals of size n. Denote the k'th order statistic by $s_{(k)}$ and its distribution by $f_k(\cdot)$. Conditional on the bidder winning the auction, the expected values of the object and the signal are $\int_1^2 E[\theta|s]f_n(s)ds$ and $\int_1^2 sf_n(s)ds$, respectively. Thus, the equilibrium bias when bidders are naive is

$$b = \int_1^2 sf_n(s)ds - \int_1^2 E[\theta|s]f_n(s)ds. \tag{16}$$

Since the bidding function is $s - b$, the auctioneer's expected revenue is $E(s_{(n-1)})$ net of the bias,

$$\int_1^2 sf_{n-1}(s)ds - b. \tag{17}$$

A Bayesian bidder bids the expected value of the object given her signal. If all bidders were Bayesian, then the winner would pay the the expected value of θ given $s_{(n-1)}$,

$$\int_1^2 E[\theta|s]f_{n-1}(s)ds. \tag{18}$$

By (16), (17), and (18), the auctioneer's revenue is higher when agents are Bayesian if and only if

$$\int_1^2 (E[\theta|s] - s)f_{n-1}(s)ds > \int_1^2 (E[\theta|s] - s)f_n(s)ds. \tag{19}$$

Condition (19) holds since $s_{(n)}$ first-order stochastically dominates $s_{(n-1)}$, and $E[\theta|s]-s$ is nonincreasing in s.

Proof of Proposition 5. Denote by $(b_t)_{t=1}^{\infty}$ the sequence of biases. Without loss of generality, we shall assume that it includes only the periods in which the signal and its corresponding realization are recorded in the DM's dataset. Let $supp(F(\cdot,\cdot)) = [\underline{s}, \overline{s}] \times [\underline{\theta}, \overline{\theta}]$. Since $b_t \in [\underline{s} - \overline{\theta}, \overline{s} - \underline{\theta}]$ for every t, it follows that if $(b_t)_{t=1}^{\infty}$ does not converge, then it traverses an interval $L^\star \subset [\underline{s} - \overline{\theta}, \overline{s} - \underline{\theta}]$ infinitely many times. In particular, it traverses (infinitely many times) an interval $[m, M] \subset L^\star$ such that $M = T(m)$ or an interval $[m, M] \subset L^\star$ such that $m = T(M)$, where $T(b)$ is the expected value of $s - \theta$ given the strategy $a^\star(s - b)$ as defined in (8). We show that the sequence cannot traverse infinitely many times an interval $[m, M]$, where $M = T(m)$. The argument for the symmetric case is analogous and, therefore, it is omitted.

Consider $b_\tau \geq M$. Let $\tilde{b}_t$ denote a sequence of biases that would be obtained if upon reaching b_τ at time τ the DM were to play $a^\star(s - m)$ each period afterwards (instead of $a^\star(s - b_t)$). Due to first-order stochastic dominance, the probability that $b_{\tau+n} \leq m$ and $b_{\tau+j} > m$ for every $j < n$ (i.e., that the sequence traverses the interval for the first time at $\tau + n$) is lower than the probability that $\tilde{b}_{\tau+n} \leq m$ and $\tilde{b}_{\tau+j} > m$ for every $j < n$. Clearly, the probability that $\tilde{b}_{\tau+n} \leq m$ and $\tilde{b}_{\tau+j} > m$ for every $j < n$ is lower than the probability that $\tilde{b}_{\tau+n} \leq m$.

Let $S_n = \sum_{i=1}^{n} (s_i - \theta_i | playing\ a^\star(s_i - m))$ and note that

(20) $$Pr\{\tilde{b}_{\tau+n} \leq m\} = Pr\{\frac{\tau b_\tau + S_n}{\tau + n} \leq m\} \leq Pr\{\frac{\tau M + S_n}{\tau + n} \leq m\} = Pr\{S_n \leq \tau(m - M) + mn\} = Pr\{S_n - nM \leq (\tau + n)(m - M)\}.$$

Since $E[s - \theta | a^\star(s - m)] = M$, the Hsu–Robbins–Erdos theorem (Hsu and Robbins, 1947; Erdos, 1949) implies that

(21) $$\sum_{n=1}^{\infty} Pr\{S_n - nM < -\epsilon n\} << \infty$$

for every $\epsilon > 0$. Let $\epsilon = M - m$. There exists $t^\star$ such that for every $\tau > t^\star$, (20) and (21) imply that

(22) $$\sum_{n=1}^{\infty} Pr\{S_n - nM < -\epsilon(\tau + n)\} < 1.$$

Hence, for a large enough τ, if $b_\tau \geq M$, then $Pr\{b_t < m\}$ for some $t > \tau$ =

$$\sum_{t=\tau}^{\infty} Pr\{b_t < m \, and \, b_j \geq m \, for \, every \, j \in \{\tau+1, ..., t-1\}\} \leq$$
$$\sum_{t=\tau}^{\infty} Pr\{\tilde{b}_t < m \, and \, b_j \geq m \, for \, every \, j \in \{\tau+1, ..., t-1\}\} \leq$$
$$\sum_{t=\tau}^{\infty} Pr\{\tilde{b}_t < m\} << 1.$$

Since this probability of traversing the interval from M to m after time τ is bounded below 1, with probability 1 the interval will be traversed a finite number of times.

We conclude that b_t must converge to some bias as t goes to infinity. Clearly, it either converges to some b for which $T(b) = b$ or to some b for which $\phi(a^\star(s-b)) = 0$ for every signal s (i.e., a corner solution).

Coarse reasoning. We now assume that the DM's prediction is

$$\hat{\theta}_{cursed} = \psi(s-b) + (1-\psi)E[\theta|s],$$

where $b = E[s - \theta|\mathcal{D}]$. We show that in an internal equilibrium, the lowest bias is increasing in ψ (the proof for the highest bias is identical). In an internal equilibrium, $T_\psi(b) = b$, where

$$T_\psi(b) := \frac{\int_{-\infty}^{\infty} f(s)\phi(a^\star(g_\psi(s))(s - E[\theta|s])ds}{\int_{-\infty}^{\infty} f(s)\phi(a^\star(g_\psi(s)))ds},$$

and

$$g_\psi(s) = \psi(s-b) + (1-\psi)E[\theta|s].$$

Note that $T_\psi(b) = b$ if and only if

$$\int_{-\infty}^{\infty} f(s)\phi(a^\star(g_\psi(s)))(s - b - E[\theta|s])ds = 0. \tag{23}$$

Assume that $T_\psi(b^\star) = b^\star$, i.e., that $b^\star$ is part of an equilibrium (specifically, the one with the lowest bias). Let $s^\star$ satisfy $s^\star - b^\star = E[\theta|s^\star]$. Note that $g_\psi(s^\star)$ is independent of ψ. Moreover, $g_\psi(s)$ is decreasing (resp., increasing) in ψ for every $s < s^\star$ (resp.,

$s > s^\star$).

Since $s - b^\star - E[\theta|s] \leq 0$ if, and only if, $s \leq s^\star$, the LHS of (23) is nonnegative for $\psi' > \psi$. That is, $T_{\psi'}(b^\star) \geq b^\star$ for $\psi' > \psi$. As in the last paragraph of the proof of Proposition 4, it is possible to show that after an increase in ψ the equality (23) is restored by increasing b.

Linear Prediction rules. We now assume that the DM considers a family of linear prediction rules $\hat{\theta}_{linear}(s) = \beta_0 + \beta_1 s$. Given a dataset $\mathcal{D}$, the parameters that minimize the quadratic loss function (i.e., the OLS estimators) satisfy

$$\beta_0 = E[\theta|\mathcal{D}] - \beta_1 E[s|\mathcal{D}] \quad \textit{and} \quad \beta_1 = \frac{cov(s, \theta|\mathcal{D})}{var(s|\mathcal{D})}.$$

This family of rules includes our baseline model, in which $\beta_1 = 1$, and multiplicative prediction rules, in which $\beta_0 = 0$.

In this more general model, whether the DM's prediction is lower than the true state (on average) depends on the functional form of $E[\theta|s]$. As an illustration, consider the investment decision setting (Application 1). The DM launches a project if and only if $\hat{\theta}(s) \geq c$. Thus, there is a cutoff signal $s^\star$ such that the DM launches a project if and only if $s \geq s^\star$ and the equilibrium dataset includes only signals higher than $s^\star$.

Suppose that $E[\theta|s]$ is convex in s. The calibrated function $\hat{\theta}(s) = \beta_0 + \beta_1 s$ intersects with $E[\theta|s]$ exactly twice in values higher than $s^\star$. To see this, note that if the two functions intersect at a single point, then rotating the linear prediction function around the intersection point would decrease the squared errors (pointwise). If the functions do not intersect, bringing them closer together by changing β_0 would clearly decrease the squared errors. Denote these intersection points by s_1 and s_2. Since $\mathcal{D}$ includes only signals higher than $s^\star$, it follows that $s_1 > s^\star$ and $s_2 > s^\star$.

Since the linear function $\hat{\theta}(s) = \beta_0 + \beta_1 s$ intersects with $E[\theta|s]$ twice above the cutoff $s^\star$ and $E[\theta|s]$ is convex, we have that $E[\theta|s^\star] > \beta_0 + \beta_1 s^\star$. This means that while the DM expects a revenue of c at the cutoff, the actual revenue is higher. Thus, a Bayesian DM would choose a lower implementation cutoff. Moreover, $E[\theta|s] > \beta_0 + \beta_1 s$ for every $s < s^\star$. Hence, $\beta_0 + \beta_1 E[s|s < s^\star] < E[\theta|s < s^\star]$. By definition, the DM's prediction is correct in the selected sample, and so $\beta_0 + \beta_1 E[s|s \geq s^\star] = E[\theta|s > s^\star]$. Combining this with the prediction out-of-sample prediction, we conclude that overall $E[\hat{\theta}] = \beta_0 + \beta_1 E[s] < E[\theta]$, as in our baseline model.

In an analogous manner, if $E[\theta|s]$ is concave, the same argument implies that $E[\hat{\theta}] = \beta_0 + \beta_1 E[s] > E[\theta]$ and that the DM's implementation cutoff is lower than the Bayesian cutoff. Thus, when the DM uses a linear prediction rule to assess the mapping from signals to states, her prediction can be either more optimistic or more pessimistic than the Bayesian prediction, depending on whether $E[\theta|s]$ is concave or convex. Consequently, her implementation cutoff can be less or more conservative relative to a Bayesian DM's cutoff depending on whether $E[\theta|s]$ is concave or convex.