

DISCUSSION PAPER SERIES

DP17329

Naive Calibration

Yair Antler and Benjamin Bachi

ORGANIZATIONAL ECONOMICS

CEPR

Naive Calibration

Yair Antler and Benjamin Bachi

Discussion Paper DP17329

Published 27 May 2022

Submitted 20 May 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Organizational Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Yair Antler and Benjamin Bachi

Naive Calibration

Abstract

We develop a model of non-Bayesian decision-making in which an agent obtains an estimate of the state of a relevant economic fundamental but does not know the joint distribution of the two. To make use of the estimate, she relies on an endogenously generated dataset that consists of previous estimates and state realizations. She attributes a systematic difference between the estimates and state realizations in her dataset to a systematic bias in the estimate and naively calibrates it. Her subsequent action affects the probability with which the estimate and the corresponding state realization will be recorded in the dataset that will be used in future decisions. We investigate the steady state of the naive calibration procedure and show 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.

JEL Classification: N/A

Keywords: Selection Bias, bounded rationality, Misspecified models

Yair Antler - yair.an@gmail.com

Tel Aviv University, Coller School of Management and CEPR

Benjamin Bachi - bbachi@econ.haifa.ac.il

University of Haifa

Naive Calibration

Yair Antler* and Benjamin Bachi†

May 19, 2022

Abstract

We develop a model of non-Bayesian decision-making in which an agent obtains an estimate of the state of a relevant economic fundamental but does not know the joint distribution of the two. To make use of the estimate, she relies on an endogenously generated dataset that consists of previous estimates and state realizations. She attributes a systematic difference between the estimates and state realizations in her dataset to a systematic bias in the estimate and naively calibrates it. Her subsequent action affects the probability with which the estimate and the corresponding state realization will be recorded in the dataset that will be used in future decisions. We investigate the steady state of the naive calibration procedure and show 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 use an economic variable to make an important decision, 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 10% higher

*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, Philippe Jehiel, Dotan Persitz, Ariel Rubinstein, and Rani Spiegler for helpful comments and suggestions. We also thank seminar audiences at the Hebrew University, NYU Abu Dhabi, and the University of Haifa for useful comments.

than actual outcomes. How would you use the estimate at your disposal in light of this finding?

A natural way to account for the apparent bias in the company's records is to choose an action as if your estimate were 10% lower. The problem with this approach is that the discrepancy in the records need not follow from an actual bias. In particular, systematic differences between estimates and ex post outcomes may emerge even when the estimates are unbiased if the data in the company's records suffers from selection bias (e.g., the records may depend on one's predecessors' decisions or the organizational memory). It is well known that when researchers fail to properly account for selection bias, they may end up reaching erroneous conclusions. But what happens when the person who reaches this erroneous conclusion is not an outside observer but rather a decision maker (DM) whose actions affect the very selection bias she is trying to account for? To address this question and others, we develop a model in which boundedly rational DMs try to account for a selection bias they obliquely 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 *An 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 $U[0, 1]$.*

Suppose that the entrepreneur takes estimates at face value and, therefore, launches only projects whose estimates are greater than 0.4. In the long run, if she examines the projects she implemented, our entrepreneur will observe a systematic discrepancy: while the average estimated return is $E[s|s \geq 0.4] = 0.7$, the actual return is, on average, $E(\theta|s \geq 0.4) = 0.5 + 0.2p$. If she resolves this discrepancy by adjusting her expectations downward by $E[s - \theta|s \geq 0.4] = 0.2(1 - p)$, then her expectation of θ would be $s - 0.2(1 - p)$, and so she would revise her strategy and launch only projects for which $s \geq 0.4 + 0.2(1 - p)$. In the long run, increasing the acceptance cutoff from $s = 0.4$ increases the discrepancy in the entrepreneur's records to $E[s - \theta|s \geq 0.4 + 0.2(1 - p)] = 0.2(1 - p) + 0.1(1 - p)^2$. 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 acceptance cutoff even further to $s = 0.4 + 0.2(1 - p) + 0.1(1 - p)^2$, and so on. In the limit of this process, the entrepreneur reaches a steady state in which she (i) observes a discrepancy

of $E[s - \theta | s \geq 0.4 \frac{2}{1+p}] = 0.4 \frac{1-p}{1+p}$ in her data, (ii) launches projects if and only if $s \geq 0.4 \frac{2}{1+p}$, and (iii) believes she uses a real cutoff of 0.4.

The entrepreneur in Example 1 sets a higher acceptance cutoff than a Bayesian entrepreneur.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 estimates are, on average, higher than the actual revenue. In an attempt to account for this discrepancy even though the estimate is unbiased (i.e., $E(\theta) = E(s)$), the entrepreneur sets an excessively high acceptance cutoff. As the iterative process in the example illustrates, setting a more conservative acceptance 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 acceptance cutoff, and so on.

The discrepancy between initial estimates and ex post realizations illustrated in Example 1 is prevalent in various contexts and, in particular, in the oil and mineral extraction industries. For instance, oil extraction projects are often selected based on their estimated reserves and, in ex post audits, these reserves are typically lower than initially estimated (Capen et al., 1971; Chen and Dyer, 2007). Similar phenomena occur in other contexts such as selection of road construction projects (Bajari et al., 2014) and defense procurement (Terasawa et al., 1989).

Our model is more general than the investment example, which makes it applicable in various contexts. In the model, we consider a DM who faces a sequence of similar decision problems. In each problem, the DM observes an independent estimate of the state of a relevant economic variable and then takes an action. We assume that the DM's actions can be ranked and that the higher the estimate, the higher the DM's optimal action. After taking an action, the DM observes the state's realization with a probability that depends on her action monotonically. The DM records all the realizations she observes and their respective estimates in a dataset that is available for future decisions. To illustrate our assumptions, consider a DM who relies on estimates of objects' values to bid in second-price IPV auctions. The higher the estimate she observes, the higher her optimal bid, and therefore the more likely she is to win the object and learn its actual value.

Since the DM's actions affect the data she observes and vice versa, we focus on the steady state of this system and refer to it as an *equilibrium*. An equilibrium consists

¹A conventional Bayesian entrepreneur would approve projects for which $E[\theta|s] \geq 0.4$, which implies a cutoff of $0.5 - \frac{0.1}{p}$ if $p > 0.2$ and 0 otherwise.

of a strategy (a mapping from estimates to actions) and a bias b , such that the bias is equal to the average difference between observed realizations and estimates, and the DM's strategy is a best response to a belief that each estimate s reflects the true state θ up to a bias b , i.e., a belief that $\theta = s - b$. We show that an equilibrium always exists. Moreover, in equilibrium, the bias is always positive and, as a result, leads to a pessimistic interpretation of the estimates as illustrated in Example 1.

We refer to the mapping from the actions taken by the DM to the probability that she observes the actual realizations ex post as a *feedback function*. The feedback function reflects basic properties of the environment in which the DM operates; for instance, different feedback functions may represent different types of organizational memory. In Example 1, the feedback function assigns a probability of 1 to implemented projects and a probability of 0 to unimplemented ones. Our model allows for various other feedback functions. We derive a tight condition that enables us to rank different feedback function in terms of the bias that they induce. Essentially, a feedback function ϕ induces a higher bias than the feedback function ϕ' for every information structure if, and only if, ϕ dominates ϕ' in the likelihood ratio sense. In Section 4 we illustrate the significance of this condition in a setting of a second-price IPV auction in which bidders rely on an estimate to choose their bids under the assumption that the actual value of the object is observed only by the winning bidder. We use the fact that the feedback depends on the number of bidders and show that the feedback-ranking condition implies that the bidders' bias is increasing in the number of bidders for any objective mapping from states to estimates.

We then study the externalities DMs impose on other naive DMs who rely on the data they generate. 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 naive calibration heuristic while the others are Bayesian in the sense that they know the joint distribution of states and estimates. Since they know the joint distribution, Bayesian agents ignore the dataset when taking an action. On the other hand, their actions select different 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.

We show that calibrating in an attempt to account for the discrepancy in the data can lead to seemingly pessimistic behavior. This is demonstrated, in Section 4,

where we apply our model to study project selection and IPV second-price auctions. This microfoundation for irrational pessimism in equilibrium is a contribution to the literature on misspecified beliefs. While there is much evidence and many models of overoptimism, the literature has not paid as much 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 effects we find are reminiscent of the winner’s curse in IPV auctions (Compte, 2002), choice-driven optimism (Van den Steen, 2004), and the optimizer’s curse (Smith and Winkler, 2006). Van den Steen and Smith and Winkler 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 different state realizations into the dataset, which requires an equilibrium approach that is absent from these papers. We discuss the differences between our model and Compte’s (2002) model of IPV procurement auctions in Section 4.

This paper contributes to the literature that investigates 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 4, we discuss the decision procedure in the latter paper and explain why it leads to opposite results from ours.

At a broader level, this paper belongs to a growing literature that studies 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; Spiegel, 2016). For excellent reviews on this topic see Jehiel (2020) and Spiegel (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, that captures Bayesian learning under a misspecified model of the world. In a Berk–Nash equilibrium the DM’s beliefs minimize the relative entropy with respect to her (misspecified) model of the world. By contrast, our solution concept imposes that observed discrepancies between

estimates and state realizations be resolved by fitting the mean of the estimate to the mean of the state.

The paper proceeds as follows. We present the model in Section 2 and analyze it in Section 3. This analysis is used in Section 4 to study second-price IPV auctions and investment decision problems. Section 5 concludes. All proofs are relegated to the appendix.

2 The Model

We start by describing the setting in which decision makers (DMs) act. Subsequently, we shall introduce our solution concept.

Information structure. The environment is composed of two random variables, namely, the state of nature θ and an estimate s , which are distributed according to a continuous bivariate density function $f(\theta, s)$ defined over \mathbb{R}^2 . We denote the marginal densities by $f(\theta)$ and $f(s)$. The estimate is unbiased ex ante in the sense that $E(\theta) = E(s)$.² We assume that higher states induce higher estimates. Specifically, the marginal distributions of estimates given states satisfy the monotone likelihood ratio property. Formally, if $s > s'$, then $\frac{f(\theta, s)}{f(\theta, s')}$ is nondecreasing in θ . Finally, we assume that $s - E[\theta|s]$ is increasing in s , namely, that the average difference between the estimate and the state is increasing in s . This property is satisfied by many prevalent information structures. In particular, the estimate can be the sum of the state and a zero-mean noise term that are drawn from log concave distributions.³

Strategies, actions, and payoffs. A strategy is a mapping $a : \mathbb{R} \rightarrow A$ from estimates to actions, where the set of actions A is a compact subset of \mathbb{R} . The DM's payoff $\pi(\theta, a)$ is a function of the state and her action. We denote the action that maximizes the payoff function given a state θ by $a^*(\theta)$ and assume that it is unique, except for, possibly, a measure zero of realizations. We assume that $a^*(\theta)$ is weakly increasing in θ and nondegenerate.

Feedback, bias, and equilibrium. To make use of the estimate at her disposal, the DM uses an infinitely large dataset that consists of estimates and state realizations from past decisions. The probability that each pair consisting of an estimate and

²While the unbiasedness assumption facilitates the exposition, with minor adjustments our results and intuitions continue to hold when the estimates are biased.

³Formally, suppose that $s = \theta + \epsilon$, where ϵ is drawn from a continuous density $h(\epsilon)$ such that $E(\epsilon) = 0$. Efron (1965) shows that log-concavity implies that $s - E[\theta|\theta + \epsilon = s]$ is increasing in s .

a state is recorded in the dataset is equal to the probability that the DM observes the realization of the state, which in turn depends on the action she takes given the estimate. Formally, the DM records each pair (s, θ) with probability $\epsilon + (1 - \epsilon)\phi(a(s))$, where $\phi : A \rightarrow [0, 1]$ is a feedback function and $a(s)$ is the action the DM takes given s . We assume that ϕ is nondegenerate and nondecreasing. Moreover, we assume that $\epsilon > 0$ in order to guarantee that the DM's beliefs are well defined and study the case in which $\epsilon > 0$ is sufficiently close to zero.

Given a strategy a , denote the average difference between states and estimates in the DM's dataset by $b(a)$. Formally,

$$(1) \quad b(a) = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} ((1 - \epsilon)\phi(a(s)) + \epsilon)f(\theta, s)[s - \theta]dsd\theta}{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} ((1 - \epsilon)\phi(a(s)) + \epsilon)f(\theta, s)dsd\theta}.$$

The DM believes that the estimate she observes reflects the true state, up to a potential bias, and she heuristically calibrates it based on the average bias in the dataset. That is, the DM believes that $\theta = s - b(a)$ and chooses optimally given that belief. Our solution concept also requires that the DM's perceived bias be consistent with the discrepancy in the data that is generated by her strategy. Formally:

Definition 1 *A strategy a and a bias b constitute an equilibrium if $a(s) = a^*(s - b)$ for every s , where $b = b(a)$ according to (1).*

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 the existence result, note that the bias b pins down the DM's strategy $a^*(s - b)$, which pins down the long-run dataset along with a new perceived bias $T(b) := b(a^*(s - b))$. The proof of Proposition 1 establishes that this process has a fixed point; i.e., there exists a bias b such that $T(b) = b$. To see this, note that

when the bias b goes to infinity, the DM essentially takes the same action regardless of the estimate's realization.⁴ This implies that $\phi(a(s))$ converges pointwise to some k_l so that in the limit the DM observes all pairs (s, θ) with the same probability, $(1 - \epsilon)k_l + \epsilon$. Since $E(\theta) = E(s)$, it follows that the induced bias $T(b)$ goes to zero when b goes to infinity. Symmetrically, $\lim_{b \rightarrow -\infty} T(b) = 0$. Since $\lim_{b \rightarrow \infty} T(b) = \lim_{b \rightarrow -\infty} T(b) = 0$, the continuity of $T(b)$ in b implies that there exists a bias b for which $T(b) = b$.

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

In equilibrium, $s - b$ is the naive DM's expectation of θ . Let us compare these beliefs to the posterior beliefs of a Bayesian DM. While the former believes that $\theta = s - b$, the latter believes that, in expectation, the state is $E[\theta|s]$. Recall that, by assumption, $s - E(\theta|s)$ is increasing in s and that 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, our DM's belief about the state is lower than a Bayesian DM's belief for low estimates and higher for high estimates. Corollary 1 summarizes this discussion.

Corollary 1 *In equilibrium, there exists an estimate \hat{s} such that a Bayesian DM's belief about the expected value of the state is weakly lower (resp., higher) than the naive DM's belief if and only if $s > \hat{s}$ (resp., $s < \hat{s}$).*

The next result shows that the equilibrium bias is nonnegative. Thus, although our DM's beliefs are higher than a Bayesian DM's beliefs for high estimates, they are lower than the Bayesian DM's beliefs on average.

Proposition 2 *In every equilibrium, the bias is nonnegative.*

To understand this result, recall that if the DM were to observe all estimates and their respective realizations, there would be no bias at all. However, the DM's dataset includes only a selected sample of such pairs. In particular, since ϕ is increasing in a and a^* is increasing in s , the dataset contains disproportionately more cases in which the estimate is high. The assumption that $s - E(\theta|s)$ is increasing 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.

⁴When b goes to infinity, $a^*(s - b)$ converges pointwise to $a_l^*(s) := \lim_{t \rightarrow -\infty} a^*(t)$ for every s .

A key aspect of our model is the selection bias in the DM's data. We now provide a definition that allows us to rank different feedback functions according to the amount of selection they induce in the data. We then use this condition to show how the amount of selection bias affects the equilibrium bias.

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 increasing in a .*

The dominant feedback function returns relatively more observations in which the DM takes high actions and fewer observations in which the DM takes low actions compared to the dominated feedback function. As an illustration, consider n DMs who participate in a second-price IPV auction and use their estimate to bid in the auction. Since bidding one's expected value is a dominant strategy, this game can be thought of as a decision problem. If a bidder observes the object's realization only when she wins the object, then the feedback function in a symmetric equilibrium induces a probability of $F(s)^{n-1}$. Thus, the feedback function when there are n bidders dominates the feedback function when there are $n - 1$ bidders in the likelihood ratio sense.

The next result establishes that if a feedback function dominates another feedback function in the likelihood ratio sense, then it induces a higher bias. Denote the minimal and maximal equilibrium biases given a feedback function ϕ by \underline{b}_ϕ and \bar{b}_ϕ , respectively.

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 $\bar{b}_\phi \geq \bar{b}_{\tilde{\phi}}$.*

The dominant feedback function ϕ puts relatively more weight on high actions. Thus, intuitively, it puts more weight on high estimates and, as a result, more weight on instances in which the estimate is high relative to the actual state realization. The more weight a feedback function puts on these instances relative to instances in which estimates 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 , a payoff function π , and an optimal strategy a^* 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 that is concentrated on that interval, a payoff

function π , and a strictly increasing function a^* , such that the result of Proposition 3 is reversed. The following corollary summarizes this discussion.

Corollary 2 *If ϕ does not dominate $\tilde{\phi}$ in the likelihood ratio sense then there exist a distribution f and a function a^* such that $\underline{b}_\phi < \underline{b}_{\tilde{\phi}}$ and $\bar{b}_\phi < \bar{b}_{\tilde{\phi}}$.*

So far, we have assumed that the dataset on which the naive DM bases her decisions is generated by actions of other naive DMs who use her same reasoning. However, it is also plausible that some of these DMs are more experienced and thus have a better knowledge of the joint distribution of estimates and states. Such DMs can apply the Bayesian machinery to make use of the estimates at their disposal. Since their actions select different estimates into the naive DM's dataset, they affect the discrepancy that the naive DM observes and, as a result, they affect her actions and the realizations they select into the dataset.

We now incorporate this idea into our model by assuming that a share $\alpha > 0$ of the DMs are Bayesians. Since Bayesian DMs are unaffected by the data, varying their share enables us to study the externalities between different 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.

For the next result, we assume that given an estimate s , a Bayesian agent plays the action $a^*(E(\theta|s))$.⁵ Let $b_B(s) = s - E(\theta|s)$ and

$$b_\alpha(a) = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} ((1 - \epsilon)[(1 - \alpha)\phi(a(s)) + \alpha\phi(a^*(E(\theta|s)))] + \epsilon)f(\theta, s)[s - \theta]dsd\theta}{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} ((1 - \epsilon)[(1 - \alpha)\phi(a(s)) + \alpha\phi(a^*(E(\theta|s)))] + \epsilon)f(\theta, s)dsd\theta}$$

represent the bias in the dataset given that a share α of the decisions are made by Bayesian DMs.

It is possible to use similar arguments to those used in the proofs of Propositions 1 and 2 to show that (i) an equilibrium exists and (ii) the bias is nonnegative for any $\alpha \in [0, 1]$. We denote the highest and lowest equilibrium biases by \bar{b}_α and \underline{b}_α , 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.

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

Proposition 4 *Both \bar{b}_α and \underline{b}_α are weakly decreasing in α .*

Relative to our naive DMs, Bayesian DMs put a lower weight on the estimate as they take their prior beliefs into account. Therefore, Bayesian DMs play lower (resp., higher) actions when the estimate is high (resp., low). Since the feedback function induces a probability that is increasing in a , it follows that actions taken by a Bayesian DM generate less (resp., more) feedback in situations in which the estimate 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.

Note that when the share of boundedly rational DMs vanishes, the bias in the data does not disappear. Indeed, $b_1 > 0$ except for the degenerate case in which $\phi(a^*(s))$ is constant for all s . The reason for this positive bias is that the selection in the dataset does not vanish when DMs become Bayesians: they still take higher actions given higher estimates, which implies that higher estimates 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 one than the bias obtained when DMs use the naive calibration procedure.

4 Applications

In this section, we apply our results to two settings: project selection and second-price IPV auctions. The main message in both applications is that our naive calibration procedure leads to conservative behavior relative to the behavior of a Bayesian DM: underbidding in auctions and rejection of marginally good projects in project selection problems.

4.1 Project Selection

An entrepreneur selects which projects to implement based on their estimated revenue. We denote the revenue by θ , and its estimate by s . Implementing a project entails a (known) cost of c . Denote a decision to implement a project by $a = 1$ and a decision to forgo it by $a = 0$. Thus, the entrepreneur's payoff is $\pi(a, \theta) = a(\theta - c)$. Hence,

she wishes to implement a project if and only if $\theta \geq c$, and so $a^*(\theta) = 1$ if $\theta \geq c$ and $a^*(\theta) = 0$ otherwise.

The entrepreneur bases her decisions on a dataset that includes estimates and actual revenues of projects selected by a sequence of entrepreneurs of which a share α are Bayesian and a share $1 - \alpha$ are naive, where $0 \leq \alpha \leq 1$. The dataset contains feedback only about implemented projects, namely,⁶ $\phi(a) = a$. Denote the entrepreneur's bias by b_α (Claim 1 shows that the equilibrium is unique and, therefore, b_α is well defined). Since she believes that the project's revenue is $s - b_\alpha$ she launches projects whose estimates are greater than $c + b_\alpha$. Hence, in equilibrium, she uses an acceptance cutoff of

$$(2) \quad s_\alpha = c + b_\alpha.$$

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

The next result shows that under a mild technical assumption the equilibrium is unique and characterizes the equilibrium bias as a function of α .

Claim 1 *If $f(s)$ is log concave, then there exists a unique equilibrium for every α . Moreover, it holds that (i) $s_\alpha \geq c$, (ii) $s_\alpha \geq s_B$, and (iii) s_α is decreasing in α .*

While parts (i) and (iii) follow almost immediately from Propositions 2 and 4, part (ii) provides new insights. This part shows that despite the fact that for sufficiently high estimates the naive entrepreneur is more optimistic than the Bayesian entrepreneur about the quality of the project (Corollary 1), *the naive entrepreneur's acceptance cutoff is weakly higher than the Bayesian entrepreneur's cutoff*. To obtain intuition for this effect, note that since $E[\theta|s = s_B] = c$, it holds that $E[s - \theta|s = s_B] = s_B - c$. Were the naive entrepreneur to use the Bayesian cutoff s_B , then her bias would be $E[s - \theta|s > s_B]$, which is greater than $E[s - \theta|s = s_B]$ since $E[s - \theta|s]$ is increasing. Thus, from the naive entrepreneur's perspective, the Bayesian entrepreneur's cutoff is too low.

The results in this section run counter to the results obtained in Jehiel's (2018) equilibrium model of project selection. In particular, while the setting in Jehiel (2018)

⁶Formally, we study the limit case in which $\epsilon \rightarrow 0$.

is similar to the one in the present section, in equilibrium, the agents use an acceptance cutoff that is too low relative to Bayesian agents. It is instructive to discuss the different modeling assumptions that lead to this difference.

In both models, the entrepreneur has access to the set of all projects that were implemented in the past and their realized revenues. In Jehiel’s model, the entrepreneur obtains a *new* signal (interpreted as an impression) about each of these projects as well as about a current project that is up for selection. Jehiel’s entrepreneur calculates the average realized revenue of past projects conditional on each signal and believes that, in expectation, the current project’s revenue will be identical to the expected revenue of projects for which she obtained the same signal. Jehiel shows that if the entrepreneurs’ signals are independent of one another, then in the equilibrium of this procedure they set an acceptance 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, that is, projects for which previous entrepreneurs received a signal above their acceptance cutoff. The expected value of these projects is higher than the expected value of unimplemented ones conditioned on the signal the entrepreneur receives. The entrepreneur ignores this selection. Since these projects are of higher quality than the pool of all projects—to use our terminology—the entrepreneur believes that her signal is downward biased and, therefore, adjusts her acceptance cutoff downward.

In both Jehiel’s model and in ours, the dataset that the entrepreneur uses to evaluate new projects suffers from selection bias. However, the datasets in both models are different. While Jehiel’s entrepreneur uses a dataset that contains relatively high states conditional on the signals, our entrepreneur uses a dataset that contains high estimates conditional on the states. Thus, while Jehiel’s entrepreneur concludes that the signals are downward biased, our entrepreneur concludes that the estimates are upward biased.

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

While we have been using the project selection terminology, the analysis in this section is relevant in other contexts as well. For example, the DM might be a physician who decides which patients to treat based on the results of a medical test and a dataset that includes test results and treatment results for previous treated patients. Alternatively, the DM might be an individual who uses a recommendation system and naively calibrates the recommendation she receives based on her actual enjoyment of a product

in previous situations in which she followed the recommendation and 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 contexts our results imply a seemingly pessimistic behavior: setting a bar too high for medical treatment and credit provision, and downgrading positive recommendations from an algorithm or even from a friend.

4.2 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 naturally extended to strategic situations in which multiple players interact. This requires extending the payoff and the feedback function, and making assumptions on how agents reason about other agents' behavior. 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, each of whom receives an estimate s_i of the value she will derive from the object, θ_i . The value and its estimate are independently drawn from $f(\theta, s)$ for each bidder.⁷ Recall that in a second-price IPV auction bidding one's value is a dominant strategy, i.e., $a^*(\theta_i) = \theta_i$. Denote bidder i 's bias by b_i . We interpret $s_i - b_i$ as agent i 's calibrated estimated value and, therefore, assume that each bidder i uses the bidding strategy $a(s_i) = s_i - b_i$. We assume that a bidder learns the true value of the object if and only if she wins the object.

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 constitutes an equilibrium at the individual level.*

In a symmetric equilibrium, all bidders use the same bidding function, receive the same feedback, and reach the same conclusion about their own bias (i.e., $b_i = b_j$ for every pair of bidders i, j). Therefore, a bidder obtains feedback if and only if her

⁷To bypass the problem of potentially negative bids, we assume that $\text{supp}(f) = [1, 2]^2$.

estimate is the highest, i.e., $\phi(a(s_i, s_{-i})) = F(s_i)^{n-1}$. Note that this feedback function is independent of the actual bidding. Hence, there is a unique bias that is consistent with our calibration procedure. The next corollary summarizes this discussion.

Corollary 3 *There exists a unique symmetric equilibrium.*

The next claim provides comparative statics with respect to the number of bidders.

Claim 2 *In a symmetric equilibrium, $a(s) < s$ and is decreasing in n .*

In a symmetric equilibrium, the bias is strictly positive. The strict inequality follows from the bids being strictly increasing in the estimates. Thus, different state realizations are observed with different frequencies, which precludes the possibility of a corner solution in which the bias is null. The comparative statics w.r.t. the number of bidders follow directly from Proposition 3. To see this, recall that the feedback probability when there are n bidders is $F(s)^{n-1}$ and this function is dominated by $F(s)^{m-1}$ for $m > n$ bidders.

As in a second-price IPV auction with Bayesian bidders, the bidder with the highest estimate wins the object. Thus, the equilibrium outcome is efficient. However, the naive calibration procedure affects the bidding strategy and, therefore, may affect how the total surplus is divided. 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. 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. To see this, recall that, by Corollary 1, for high (resp., low) estimates a naive bidder bids higher (resp., lower) than a Bayesian bidder.

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

This result follows from comparing the expected bid given the second highest estimate. Denote the highest and second highest estimates 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 estimate and the expected value conditional on receiving the highest estimate (recall that bidders obtain feedback only if they win the auction). Thus, a naive winner pays, on average, $E(s_{(n-1)}) - (E(s_{(n)}) - E(\theta|s_{(n)}))$. When agents are Bayesian, the bidder pays the expected value of the object given the second

highest estimate, $E(\theta|s_{(n-1)})$. 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.

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, conditional on winning, the estimated cost is likely to be lower than the actual cost. Compte (and later Compte and Postlewaite, 2019) assumes that bidders use a specific bidding function: they choose a fixed markup and add it to their estimated cost to correct for the selection bias. In equilibrium, this markup is positive and increasing in the number of bidders (as in our model, the larger the number of bidders, the stronger the selection).

It is straightforward 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). To gain intuition for the larger adjustment in our model, note that our bidders take into account the difference between the actual value and the estimated value conditional on winning, whereas the bidders in Compte (2002) correct for the difference between the actual value and the estimated value conditional on the pivotal realization. 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.

5 Concluding Remarks

DMs are often provided with estimates or recommendations that are generated by analysts or AI-based algorithms. If ex post audits indicate that initial estimates were too high, DMs may become skeptical and discount these recommendations. This skepticism may result in aborting profitable projects and, in general, in behavior that appears to be excessively risk averse. Our analysis suggests that this conservative behavior may widen the gap between estimates and ex post realizations even further, which, in turn, may reinforce DMs’ skepticism.

A classic question in the behavioral economics literature is whether higher degrees of awareness or sophistication result in better outcomes. In recent years, it has been shown

that in strategic interactions, more sophisticated players may obtain lower payoffs.⁸ Our analysis suggests that DMs who trust analysts’ recommendations and take them at face value obtain higher payoffs than DMs who rely on ex post audits to calibrate their estimates. This can be illustrated using the project selection application. If the expected value of the project is higher than the cost, then the acceptance cutoff set by a DM who takes estimates at face value lies between the excessively high acceptance cutoff set by a DM who uses our naive calibration procedure and the acceptance cutoff that Bayesian a DM would use. Thus, the DM who uses our naive calibration procedure rejects more (profitable) projects than a DM who takes estimates at face value. We conclude that a greater degree of sophistication may actually lead to outcomes that are systematically worse even when there is no strategic interaction.

In many situations DMs’ estimates are provided by a strategic agent. For instance, a buyer may receive noisy information about the suitability or quality of different products advertised by a strategic seller. Such a seller might have an incentive to exaggerate the merits of products or add noise to the transmitted information. However, a DM who uses the naive calibration procedure not only corrects for the actual bias, but also interprets the noise as an additional systematic upward bias and corrects for that. 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. Hence, the naive calibration procedure may protect DMs from strategic obfuscation and misinformation.

References

- BAGNOLI, M. AND T. BERGSTROM (2005): “Log-concave probability and its applications,” *Economic theory*, 26, 445–469.
- BAJARI, P., S. HOUGHTON, AND S. TADELIS (2014): “Bidding for incomplete contracts: An empirical analysis of adaptation costs,” *American Economic Review*, 104, 1288–1319.
- BARRON, K., S. HUCK, AND P. JEHIEL (2019): “Everyday econometricians: Selection neglect and overoptimism when learning from others,” Unpublished Manuscript.
- CAPEN, E. C., R. V. CLAPP, AND W. M. CAMPBELL (1971): “Competitive bidding in high-risk situations,” *Journal of petroleum technology*, 23, 641–653.

⁸See, e.g., Ettinger and Jehiel (2010) and Eyster and Piccione (2013).

- CHEN, M. AND J. DYER (2007): “Inevitable disappointment in projects selected based on forecasts,” in *Hydrocarbon Economics and Evaluation Symposium*, OnePetro.
- COMPTE, O. (2002): “The winner’s curse with independent private values,” *Working Paper*.
- COMPTE, O. AND A. POSTLEWAITE (2019): *Ignorance and uncertainty*, 61, Cambridge University Press.
- EFRON, B. (1965): “Increasing properties of Polya frequency function,” *The Annals of Mathematical Statistics*, 272–279.
- ESPONDA, I. (2008): “Behavioral equilibrium in economies with adverse selection,” *American Economic Review*, 98, 1269–91.
- 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.
- 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, 1582–97.
- (2020): “Analogy-based expectation equilibrium and related concepts: Theory, applications, and beyond,” Unpublished Manuscript.

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.

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.

TERASAWA, K., J. QUIRK, AND K. WOMER (1989): “Turbulence, cost escalation and capital intensity bias in defense contracting,” in *Cost Analysis Applications of Economics and Operations Research*, Springer, 97–116.

VAN DEN STEEN, E. (2004): “Rational overoptimism (and other biases),” *American Economic Review*, 94, 1141–1151.

Appendix: Proofs

Proof of Proposition 1. To establish equilibrium existence, we show that there exists a bias b^* such that $T(b^*) = b^*$. Since

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \epsilon f(\theta, s) [s - \theta] ds d\theta = 0,$$

we can incorporate $a(s) = a^*(s - b)$ into the RHS of (1) and denote $\delta = \frac{\epsilon}{1-\epsilon}$ to obtain the continuous operator

$$(3) \quad T(b) = \frac{\int_{-\infty}^{\infty} f(s) (\phi(a^*(s - b))) [s - E(\theta|s)] ds}{\int_{-\infty}^{\infty} f(s) \phi(a^*(s - b)) ds + \delta}.$$

Since $\lim_{b \rightarrow \infty} \phi(a^*(s - b)) \in [0, 1]$ for every s , it holds that

$$\lim_{b \rightarrow \infty} \int_{-\infty}^{\infty} f(s) (\phi(a^*(s - b))) [s - E(\theta|s)] ds = \underline{\phi} \int_{-\infty}^{\infty} f(s) [s - E(\theta|s)] ds = 0$$

for some $\underline{\phi} \in [0, 1]$. Similarly, since $\lim_{b \rightarrow -\infty} \phi(a^*(s - b)) \in [0, 1]$ for every s , it holds

that

$$\lim_{b \rightarrow -\infty} \int_{-\infty}^{\infty} f(s)(\phi(a^*(s-b)))[s - E(\theta|s)]ds = \bar{\phi} \int_{-\infty}^{\infty} f(s)[s - E(\theta|s)]ds = 0$$

for some $\bar{\phi} \in [0, 1]$. Note that the denominator of (3) is positive at both limits. Hence, $\lim_{b \rightarrow \infty} T(b) = \lim_{b \rightarrow -\infty} T(b) = 0$. Since $T(b)$ is continuous, there exists b^* for which $T(b^*) = b^*$.

Proof of Proposition 2. Fix an arbitrary b . Let s^* denote an arbitrary estimate that satisfies $s^* - E[\theta|s^*] = 0$. Thus, $s \leq E[\theta|s]$ for $s < s^*$ and $s \geq E[\theta|s]$ for $s > s^*$. Since a^* is increasing in s , and ϕ is increasing in a , we obtain that

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

and

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

The sum of the RHS of the two inequalities is 0 as $E(\theta) = E(s)$. Thus, the numerator of the RHS of (3), which is equal the sum of the LHS of the two inequalities, is non-negative. Furthermore, its denominator is positive. Hence, $T(b) \geq 0$ for any b .

Proof of Proposition 3. To prove this result, we show that $T_{\phi}(b) \geq T_{\tilde{\phi}}(b)$ for any b , where T_{ϕ} (resp., $T_{\tilde{\phi}}$) denotes the operator T when the feedback function is ϕ (resp., $\tilde{\phi}$). For a sufficiently small $\epsilon > 0$, this is equivalent to

$$(4) \quad \frac{\int_{-\infty}^{\infty} f(s)\phi(a^*(s-b))[s - E(\theta|s)]ds}{\int_{-\infty}^{\infty} f(s)\phi(a^*(s-b))ds} \geq \frac{\int_{-\infty}^{\infty} f(s)\tilde{\phi}(a^*(s-b))[s - E(\theta|s)]ds}{\int_{-\infty}^{\infty} f(s)\tilde{\phi}(a^*(s-b))ds}.$$

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

$$(5) \quad \int_{-\infty}^{\infty} f(s)\phi(a^*(s-b))ds = \int_{-\infty}^{\infty} f(s)\tilde{\phi}(a^*(s-b))ds.$$

Since $\frac{\phi(a^*(s-b))}{\tilde{\phi}(a^*(s-b))}$ is increasing in s , (5) implies that there exists s^* such that $\phi(a^*(s^*-b)) = \tilde{\phi}(a^*(s^*-b))$. Thus, $\phi(a^*(s^*-b)) \geq \tilde{\phi}(a^*(s^*-b))$ for $s \geq s^*$ and the inequality is reversed

for $s \leq s^*$. Using the normalization in (5), we can write (4) as

$$(6) \quad \int_{-\infty}^{s^*} f(s) \left(\phi(a^*(s-b)) - \tilde{\phi}(a^*(s-b)) \right) [s - E(\theta|s)] ds + \int_{s^*}^{\infty} f(s) \left(\phi(a^*(s-b)) - \tilde{\phi}(a^*(s-b)) \right) [s - E(\theta|s)] ds \geq 0.$$

Since $s - E(\theta|s)$ is increasing in s , the LHS of (6) is greater than

$$(7) \quad \int_{-\infty}^{s^*} f(s) \left(\phi(a^*(s-b)) - \tilde{\phi}(a^*(s-b)) \right) [s^* - E(\theta|s^*)] ds + \int_{s^*}^{\infty} f(s) \left(\phi(a^*(s-b)) - \tilde{\phi}(a^*(s-b)) \right) [s^* - E(\theta|s^*)] ds,$$

which, by (5), is equal to zero.

Proof of Proposition 4. In a similar manner to (3), $T_\alpha(b_\alpha)$ can be written as

$$(8) \quad T_\alpha(b_\alpha) := \frac{\int_{-\infty}^{\infty} f(s) [(1-\alpha)\phi(a^*(s-b_\alpha)) + \alpha\phi(a^*(s-b_B(s)))] E[s-\theta|s] ds}{\int_{-\infty}^{\infty} f(s) [(1-\alpha)\phi(a^*(s-b_\alpha)) + \alpha\phi(a^*(s-b_B(s)))] ds + \delta}.$$

We first show that (8) is decreasing in α when $T_\alpha(b_\alpha) = b_\alpha$. The derivative of (8) with respect to α is negative if and only if

$$(9) \quad \int_{-\infty}^{\infty} f(s) [\phi(a^*(s-b_B(s))) - \phi(a^*(s-b_\alpha))] (E[s-\theta|s] - T_\alpha(b_\alpha)) ds \leq 0.$$

Replacing $T_\alpha(b_\alpha)$ by b_α yields

$$(10) \quad \int_{-\infty}^{\infty} f(s) [\phi(a^*(s-b_B(s))) - \phi(a^*(s-b_\alpha))] (b_B(s) - b_\alpha) ds \leq 0.$$

Since $b_B(s) = E[s-\theta|s]$ is increasing in s and b_α is independent of s , we obtain that either $\phi(a^*(s-b_B(s))) - \phi(a^*(s-b_\alpha)) \leq 0$ or $b_B(s) - b_\alpha \leq 0$, which implies (10).

Note that $\lim_{b_\alpha \rightarrow \infty} T_\alpha(b_\alpha) << \infty$ and $\lim_{b_\alpha \rightarrow -\infty} T_\alpha(b_\alpha) >> -\infty$. As a result, $T_\alpha(b_\alpha)$ crosses b_α from above at \underline{b}_α and \bar{b}_α . Inequality (10) implies that when α is increased these two crossing points must occur at a weakly lower b_α .

To complete the proof, we need to make sure that an increase in α does not add a new point $b_\alpha > \bar{b}_\alpha$ at which $T_\alpha(b_\alpha) = b_\alpha$. First, assume that (10) holds with strict

inequality for every α and b_α that satisfy $T_\alpha(b_\alpha) = b_\alpha$. By continuity, (9) holds with strict inequality in a close neighborhood of α, b_α . Thus, an increase in α cannot add a new crossing point above \bar{b}_α .

Suppose that (10) holds as an equality for some $\hat{\alpha}$ and $b_{\hat{\alpha}}$ that satisfy $T_{\hat{\alpha}}(b_{\hat{\alpha}}) = b_{\hat{\alpha}}$. It follows that $\phi(a^*(s - b_B(s)) - \phi(a^*(s - b_{\hat{\alpha}})) = 0$ for almost every $s \in (-\infty, \infty)$. Fixing $b_{\hat{\alpha}}$, we obtain that the expression $\phi(a^*(s - b_B(s)) - \phi(a^*(s - b_{\hat{\alpha}}))$ is independent of α . Hence, (9) holds with equality for α in the neighborhood of $\hat{\alpha}$. We conclude that an increase in α cannot generate new points at which $T_\alpha(b) = b$ above \bar{b}_α .

Proof of Claim 1. Since the entrepreneur chooses $a = 1$ for projects with an estimated return above s_α and $a = 0$ otherwise, $\phi(a) = a$ implies that $b_\alpha = E[s|s \geq s_\alpha] - E[\theta|s \geq s_\alpha]$. Hence, (2) can be written as

$$(11) \quad s_\alpha - E[s|s \geq s_\alpha] = c - E[\theta|s \geq s_\alpha].$$

Since $E[\theta|s]$ is increasing in s , the RHS of (11) is decreasing in s_α . The log-concavity of the marginal density $f(s)$ implies that the LHS is increasing in s_α (Bagnoli and Bergstrom, 2005, Lemma 2). Therefore, the equilibrium is unique.

Note that (i) holds if $b_\alpha \geq 0$. Since $E(s) = E(\theta)$ and $s - E[\theta|s]$ is increasing in s , it holds that $b_\alpha = E[s - \theta|s \geq s_\alpha] \geq 0$ for any s_α .

To prove (ii), observe that were a naive entrepreneur to use a cutoff s_B , she would perceive a bias of

$$E[s|s \geq s_B] - E[\theta|s \geq s_B] \geq s_B - E[\theta|s_B],$$

where the inequality stems from the assumption that $s - E[\theta|s]$ is increasing in s . Note that $E[\theta|s_B] = c$ and so (11) becomes

$$s_B - E[s|s \geq s_B] \leq c - E[\theta|s \geq s_B].$$

Single crossing implies that for (11) to hold, it must be that the naive entrepreneur's cutoff satisfies $s_\alpha \geq s_B$.

Finally, note that (iii) is implied directly by Proposition 4 and Condition (2).

Proof of Claim 2 and Corollary 3. If agents' strategies are symmetric, we can

write the bias as

$$(12) \quad b = \frac{\int_1^2 f(s)F(s)^{n-1}[s - E(\theta|s)]ds}{\int_1^2 f(s)F(s)^{n-1}ds + \delta}.$$

As the RHS of (12) is independent of b , it has a unique solution. By Proposition 2, $b \geq 0$ and, therefore, $a(s) \leq s$.

When the number of bidders is n , the feedback probability given an estimate s is $\phi_n(a^*(s - b_n)) = F(s)^{n-1}$. Thus, $\frac{\phi_{n+1}(a^*(s - b_{n+1}))}{\phi_n(a^*(s - b_n))} = F(s)$ is increasing in s and, by Proposition 3, the proof is complete.

Proof of Claim 3. Let f_k denote the density function of the distribution of the k -th-order statistic in a sample of n independent draws from $f(s)$, $s_{(k)}$. Conditional on winning the auction, the expected values of the object and the estimate 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

$$(13) \quad b = \int_1^2 sf_n(s)ds - \int_1^2 E[\theta|s]f_n(s)ds.$$

Since players subtract the bias from their estimates when bidding, and the winner pays the second highest bid, the auctioneer's expected revenue is $E(s_{(n-1)})$ net of the bias,

$$(14) \quad \int_1^2 sf_{n-1}(s)ds - b.$$

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

$$(15) \quad \int_1^2 E[\theta|s]f_{n-1}(s)ds.$$

By (13), (14), and (15), the auctioneer's revenue is higher when agents are Bayesian if and only if

$$(16) \quad \int_1^2 (E[\theta|s] - s)f_{n-1}(s)ds > \int_1^2 (E[\theta|s] - s)f_n(s)ds.$$

Condition (16) holds since $s_{(n)}$ first-order stochastically dominates $s_{(n-1)}$, and $E[\theta|s] - s$ is decreasing in s by assumption.