Naive Calibration

Yair Antler[*] and Benjamin Bachi[‡]

Preliminary and incomplete (please do not circulate)

February 8, 2022

Abstract

We study a model in which non-Bayesian decision makers (DMs) rely on noisy estimates of a payoff-relevant state. They have access to an endogenously created data set that includes a subset of previous estimates and state realizations. The DMs attribute systematic differences (if there are any) between the estimates and the state realizations to a systematic bias in the estimate and naively correct for it by subtracting the perceived bias in their data set from new estimates and taking the calibrated estimates at face value. We study this natural heuristic calibration procedure and show that, in equilibrium, it leads to over-correction downward. We explore the implications of this procedure for various contexts such as project approval and second-price IPV auctions.

1 Introduction

Conventional economic models typically assume that agents are impeccable Bayesians, i.e., that they have well-defined prior beliefs and that, when they receive a signal about the state of a relevant economic variable, they use the Bayesian machinery to update their prior beliefs. However, in reality, there are many situations in which agents do not have complete knowledge of the joint distribution of signals and states or lack

[*]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 Philippe Jehiel, 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.

the mental resources that are needed to perform such calculations. In such situations agents may use their information in alternative ways.

Consider an agent who faces a decision problem and has a signal of a relevant economic variable at her disposal but does not know the joint distribution of states and signals. If the agent has faced similar problems in the past, then it would be natural for her to try and deduce the relation between the signal and the state from the empirical distribution of past states and signals. However, using past data may have two limitations. First, in many situations, the availability of past data depends on past choices. For example, a credit officer is likely to observe the rate of return of only successful past applications. Second, the past information may be coarse. In the above example, the credit officer may consider the rate of return of all successful applications rather than the return per loan and the specific context in which it was approved.

In this paper, we develop a model of decision-making in which boundedly rational agents use endogenously generated partial information about past signals and realizations to update their beliefs about the current state. We refer to the updating procedure as *naive calibration.* To illustrate our calibration procedure and its implications, consider the following 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 an estimate of their returns. Implementing a project costs* 0.4 *and yields a revenue of* $\theta \sim U[0,1]$. *The estimate of the project's return,* s, *equals* θ *with probability* p. *Otherwise, it is an independent draw from* $U[0,1]$.

Let us start by assuming that the entrepreneur takes estimates at face value, i.e., launches only projects whose estimates are greater than 0.4. *In the long-run, our entrepreneur will observe a systematic discrepancy: while the average estimate regarding projects that are launched is* $E[s|s \geq 0.4] = 0.7$, *they yield, on average,* $E(\theta|s \geq 0.4) = 0.5 + 0.2p$. *According to our calibration procedure, the entrepreneur resolves the discrepancy by adjusting her expectations downward by* $E[s - \theta|s \geq 0.4] = 0.2(1-p)$. *Her calibrated estimate of* θ *given the estimate* s *is therefore* $s - 0.2(1-p)$, *and so she launches only projects for which* $s \geq 0.4 + 0.2(1-p)$. *Note that increasing the cutoff increases the long run bias in the observed data 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 bias increases the entrepreneur's cutoff to* $s = 0.4 + 0.2(1-p) + 0.1(1-p)^2$, *and so on. In the limit of this* *process, the entrepreneur (i) observes a bias of* $E[s-\theta|s \geq 0.4\frac{2}{1+p}] = 0.4\frac{1-p}{1+p}$*, (ii) believes she uses a cutoff of* 0.4*, and (iii) launches projects if and only if* $s \geq 0.4\frac{2}{1+p}$*. Note that our entrepreneur sets a higher acceptance cutoff than the conventional Bayesian entrepreneur for any*[1] $p < 1$.

The entrepreneur in Example 1 calibrates the signals at her disposal by fitting the means of the estimates to the mean of the states' realizations in her data set. Note that although $E(s) = E(\theta)$ the entrepreneur's data set includes an apparent bias as, on average, estimated returns are higher than actual returns. What leads to the apparent bias in the entrepreneur's data set is the selection criterion: she has data only on projects whose estimates are relatively high as projects whose estimated returns are low are aborted. The calibration procedure leads the entrepreneur to adopt a more conservative acceptance strategy and raise her project acceptance threshold. In turn, the higher threshold exacerbates the apparent bias in the data and leads to an even more conservative acceptance strategy. In this paper, we study the equilibrium of such processes.

Similar investment problems are prevalent in many 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 (Chen and Dyer, 2007). This may lead decision makers and outside observers to wonder why profits are not as high as estimated and try and correct for this bias by calibrating or adjusting future estimates downward. In turn, the calibration process may affect future decisions and the project realizations that will be observed in the future, and so on.

In our model, we go beyond the above-mentioned investment problem and consider a DM who chooses an action (not necessarily binary) based on a signal[2] about the state of a relevant economic variable. We make relatively mild assumptions about the objective mapping from states to signals that hold, for example, when the signals are equal to the sum of the state and a zero-mean noise term. We also assume that the DM's action set is ordered in the sense that the higher the state, the higher the DM's optimal action. As for the DM's data set, we assume that, after she takes an action, she observes the state's realization with a probability that depends on her action monotonically. For example, consider a DM who uses an estimate to bid in a

[1] A conventional Bayesian DM 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.

[2] We use the words "signal" and "estimate" interchangeably throughout the paper.

second-price IPV auction. The higher the estimate, the higher the optimal bid, and the higher the DM's bid, the more likely she is to win the object and learn its actual value.

Since the DM's actions affect the data she observes, they also influence her future actions and the data she will observe in the future. We study the steady state of this system and refer to it as *equilibrium*.[3] An equilibrium consists of a strategy, a bias, and a data set such that (i) the dataset is consistent with the DM's strategy, (ii) the bias is equal to the average difference in the dataset between realizations and estimates, and (iii) the DM's strategy is optimal given the bias.

We show that, in equilibrium, the DM's data set includes a systematic discrepancy: on average, signals are higher than actual realizations. As a result, the DM believes that the signal is upward biased even in instances in which ex ante it is not. As in Example 1, this stems from the fact that the DM's data set is selected. Since the higher the estimate the higher the action taken by the DM, and the higher the action the higher the probability that the realization is observed, the rate at which the DM observes high estimates is disproportionately high. As a result, the average difference between states and realizations in the DM's data set, $E[s - \theta|s]$, is positive.

It is instructive to compare the expectation of our naive DM to those of a Bayesian DM. First, when the estimates are unbiased, the Bayesian DM's expectations are on average equal to $E(\theta) = E(s)$ by the law of iterated expectations while, as we noted above, our DM's downward correction leads her to holding lower expectations on average. Second, when a Bayesian DM makes a prediction, she puts some weight on her prior beliefs and some on the estimate and therefore some of the estimates are updated upward and some downward. By comparison, our naive DM does not have prior beliefs and, therefore, puts too much weight on the estimate: she updates all estimates downward by the same magnitude. This leads to expectations that are too high (resp., low) relative to a Bayesian DM given high (resp., low) estimates.

We explore how the DM's bias changes with the feedback structure, namely, with the mapping from actions to frequencies with which the DM observes the actual realizations. We derive a necessary and sufficient condition that enables us to rank different feedback structures 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. We illustrate

[3] This approach is standard in situations where agents fit a naive model to endogenously selected data; see, e.g., Esponda (2008), Esponda and Pouzo (2016), Spiegler (2016), and Jehiel (2018).

the importance of this condition in a setting of a second-price private-value auction in which bidders rely on an estimate to choose their bids.[4] We assume that the actual value of the object is observed only by the winning bidder. Our feedback-ranking condition essentially implies that the bidders' bias is increasing in the number of bidders for any objective mapping from states to estimates.

We study how the presence of Bayesian DMs (i.e., DMs with a correct prior belief and knowledge of the joint distribution of states and signals) affects our naive DM's calibration. Specifically, we assume that a fraction α of DMs are Bayesian and that a fraction $1 - \alpha$ use our naive calibration heuristic. We show that the naive DM's bias is decreasing in α; that is, the presence of Bayesian DMs mitigates our DM's bias. This effect stems from the fact that in the presence of low estimates, Bayesian agents take, on average, higher actions than naive agents, which results in a less "selected" dataset. Interestingly, our DM's bias does not vanish even when the proportion of Bayesian DMs approaches 1 as their decisions still result in a selected data set.

We apply our model and use our comparative statics results in various decision problems and dominance-solvable games. The main message in all of the applications is that naive calibration leads to conservative behavior, namely, rejection of marginally good projects, underbidding, and underinvestment.

The effects we find in our model are reminiscent of the optimizer's curse (Smith and Winkler, 2006) and the choice-driven optimism model of Van Der Steen (2004). In these papers, individuals sometimes overestimate and sometimes underestimate the benefits of each action and try to select the optimal one. As a result, they are more likely to select actions whose merits they overestimated. They will therefore tend to be overoptimistic ex ante (as in Van Der Steen's model) and disappointed ex post (as in Smith and Winkler's model). Loosely speaking, in our model, the ex-post disappointment leads the agents to calibrate the signals at their disposal, which, in turn, exacerbates this effect.

Our paper's contribution is twofold. First, we contribute to the literature on misspecified beliefs as we provide micro foundations for irrational pessimism in equilibrium (e.g., holding pessimistic beliefs about the project's quality in the setting of Example 1). Just like optimism, pessimism may lead individuals to erroneous conclusions and suboptimal choices that can have negative consequences. However, while there is much evidence and many models of overconfidence, the literature has not paid as much at-

[4] This result can be extended to any type of private-value auction in which, in a symmetric equilibrium, the bidder with the highest value wins the object.

tention to irrational pessimism, which is an equally important topic.

Second, we contribute to the literature that investigates naive learning from endogenously selected data. Esponda and Vespa (2018) and Barron et al. (2021) 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 micro foundations of overoptimism when there is positive selection. Unlike in these models, we focus on a discrepancy between estimates and states that is inherent when there is endogenous selection. We show that calibration that tries to resolve this discrepancy may lead to seemingly pessimistic behavior.

At a broader level, this paper belongs to a growing literature that studies decision making and strategic interaction when agents' models of the world are misspecified (Piccione and Rubinstein, 2003; Eyster and Rabin, 2005; Jehiel, 2005; Esponda, 2008; Esponda and Pouzo, 2016; Spiegler, 2016). For excellent reviews on this topic see Jehiel (2020) and Spiegler (2020). The modeling 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 from endogenously created feedback and a misspecified model of the world. A decision maker's strategy constitutes a Berk–Nash equilibrium if it is (i) optimal given the DM's beliefs and (ii) the DM holds beliefs that minimize the relative entropy with respect to her prior (misspecified) belief and the feedback she obtains given her strategy. Our solution concept imposes the first requirement (which is standard in economics) and imposes the additional requirement that any systematic discrepancies given the DM's feedback be resolved by fitting the mean of the estimate to the mean of the state. The latter requirement is equivalent to an assumption that the DM minimizes mean square error with respect to some (potentially misspecified) prior belief about the relation between the estimate and the state. In addition to this difference, our approach is less general, which allows us to provide stronger results in specific contexts.

2 The Model

We now describe the ingredients of the environment in which decisions 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 $\theta \in \mathbb{R}$ and an estimate $s \in \mathbb{R}$, which are distributed according to a continuous bivariate density function $f(\theta, s)$ defined over $\mathbb{R}^2$. We use $f(\theta)$ and $f(s)$ to denote the marginal densities. The estimate is unbiased ex ante in the sense that $E(\theta) = E(s)$.[5] We assume that higher states induce higher estimates. Specifically, the marginal distributions of estimates given states satisfy the monotone likelihood ratio property (MLRP). 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.[6]

Strategies. A strategy is a mapping $a : [\underline{s}, \overline{s}] \rightarrow A$ 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 $a^{\star}(\theta) = argmax_a \pi(\theta, a)$ and assume that it is unique, except for a measure zero of realizations. We assume that $a^{\star}(\theta)$ is weakly increasing in θ and non-degenerate.

Feedback, bias, and equilibrium. To utilize the signal at her disposal, the DM uses an infinitely large data set that includes past signals and states. The probability that each pair consisting of an estimate and a state is recorded in the data set 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 non-degenerate and weakly increasing in the DM's action. Thus, estimates for which the DM takes higher actions receive relatively large weight in her data set. We assume that $\epsilon > 0$ to guarantee that the DM's beliefs are well defined regardless of her behavior and study the case in which $\epsilon > 0$ is sufficiently close to zero.

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

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

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

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

We now provide the formal definition of our solution concept, which has two parts. The first describes the DM's beliefs given the signal and the average bias she observes. Specifically, we assume that the DM believes that the signal she observes reflects the true state, up to a potential bias, and she heuristically calibrates it based on the average bias in the data set. That is, the DM believes that $\theta = s - b(a)$. The second part requires that the DM's perceived bias be consistent with the data that is generated by her behavior. Thus, the solution concept captures the steady state of the DM's behavior and reasoning. Formally:

Definition 1 *A strategy a and a bias b form an equilibrium if:*

1. *The strategy a is optimal given b, that is, $a(s) = a^\star(s - b)$.*
2. *The bias b satisfies $b = b(a)$ according to* (1).

3 Analysis

We start the analysis by showing that an equilibrium exists. We then show that the equilibrium bias is always positive and compare our DM's behavior to the behavior of Bayesian DMs. After this comparison, we study how the DM's bias changes with the model's primitives.

The first result establishes the existence of an equilibrium.

Proposition 1 *An equilibrium exists.*

To gain intuition for the existence result, note that the bias b pins down the agent's strategy $a^\star(s - b)$, which, in turn, affects the data set at the agent's disposal, which induces a new perceived bias $T(a^\star(s - b))$. The proof of Proposition 1 establishes that this process has a fixed point; i.e., there exists a bias b such that $T(a^\star(s - b)) = b$. To see this, note that when the bias b goes to infinity, $a^\star(s - b)$ goes to $\inf(A)$ for every s, in which case the DM essentially takes the same action regardless of the signal's realization. This implies that in the limit the DM observes all pairs (s, θ) with the same probability, $(1 - \epsilon)\phi(inf(A)) + \epsilon$. Since $E(\theta) = E(s)$, it follows that the induced bias $T(a^\star(s - b))$ goes to zero. Symmetrically, when the bias b goes to minus infinity, $a^\star(s - b)$ goes to $sup(A)$. It follows that the DM takes the same action regardless of the signal's realization, which implies that $lim_{b \to -\infty} T(a^\star(s - b)) = 0$.

Since $lim_{b\to\infty}T(a^\star(s-b)) = lim_{b\to-\infty}T(a^\star(s-b)) = 0$, the continuity of $T(a^\star(s-b))$ in b implies that there exists a bias b for which $T(a^\star(s-b)) = b$.

Note that the equilibrium in our model need not be unique and, therefore, the equilibrium bias need not be unique either.[7] We thus focus on the minimal and maximal biases that are consistent with an equilibrium in our model. We denote the maximal equilibrium bias by $\overline{b}$ and the minimal equilibrium bias by $\underline{b}$.

After establishing that an equilibrium exists, we study the equilibrium bias, its direction, and the factors that affect its magnitude. We start by showing that the equilibrium bias must be nonnegative.

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

To understand this result, recall that the estimates are ex ante unbiased ($E(\theta) = E(s)$) and so, if the DM were to observe all estimates and their respective realizations, there would be no bias at all. However, the DM's data set includes only a selected sample of such pairs. In particular, since ϕ is increasing in a and a^* is increasing in s, the data set contains disproportionally more cases in which the estimate is high. The assumption that $s - E(\theta|s)$ is increasing in s implies that the data set also contains disproportionally more cases in which the bias is high, which, in turn, leads to a perceived upward bias.

Let us interpret $s-b$ as the DM's belief about θ and compare her equilibrium belief 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. Note that b is a weighted average of $s - E[\theta|s]$ and therefore $b \in (inf(s - E[\theta|s]), sup(s - E[\theta|s]))$. We conclude that our DM's belief about the state is lower than a Bayesian DM's belief given low estimates and may be higher for high estimates. Intuitively, our DM does not have a prior and therefore she puts too much weight on the signal. This results in beliefs that are too high given high signals and beliefs that are too low given low signals. The next corollary summarizes this argument.

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

[7] We wish to point out that in the subsequent applications section the equilibrium is unique.

A central aspect of our model is the selection in the DM's data. The more the DM is exposed to high estimates relative to her exposure to low estimates, the greater is the positive selection in her data set and, as a result, the greater is her bias. Since higher estimates induce higher actions, a feedback function that provides relatively more feedback when the agent takes a high action will generate a higher bias. We now provide a formal definition that captures this idea and allows us to rank different feedback functions.

Definition 2 *We say that 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 took high actions and fewer observations in which the DM took low actions compared to the dominated function. As an illustration, consider n DMs who participate in a second-price IPV auction and use their signal 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 is $F(s)^{n-1}$. 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.

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 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 $\tilde{\phi}$ in the likelihood ratio sense, then there exist a payoff function π, a corresponding function $a^\star$, and a distribution f such that $\tilde{\phi}$ induces a higher bias than ϕ. 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 both a distribution f that is concentrated on that interval, and a strictly increasing function $a^\star$, 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 $a^\star$ such that $\underline{b}_\phi < \underline{b}_{\tilde{\phi}}$ and $\overline{b}_\phi < \overline{b}_{\tilde{\phi}}$.*

So far, we have assumed that the data set on which the naive DM bases her decisions is generated by actions of other naive DMs. However, it is also plausible that other DMs are more experienced and thus have a better knowledge of the joint distribution of estimates and states, and so they can apply the Bayesian machinery to utilize the estimates at their disposal. Since their actions select different estimates into the DM's data set, they indirectly affect her actions via the discrepancy she observes.

We now incorporate this idea into our model by assuming that a share $\alpha > 0$ of the decision makers are Bayesians and explore the implications for the DM's equilibrium bias and whether the presence of these agents can bring the DM's behavior closer to Bayesian behavior. Note that, given an estimate s, a Bayesian agent plays the action $a^\star(E(\theta|s))$. Denote $b_B(s) = s - E(\theta|s)$ and let

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

represent the bias in the DM's data set.

It is possible to use a similar argument to the one used in Propositions 1 and 2 to show that (i) an equilibrium exists and (ii) that the bias is nonnegative for any $\alpha \in [0,1]$. We 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. In fact, the more Bayesian DMs there are, the lower the bias is.

Proposition 4 *Both $\overline{b}_\alpha$ and $\underline{b}_\alpha$ are weakly decreasing in α.*

Relative to our naive DM, 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). Because of the monotonicity of the feedback function ϕ in a, this implies that actions taken by a Bayesian DM generate less (resp., more) feedback in situations in which the estimate is high (resp., low) and, as a result, the average bias in the data set is lower when the data set contains more actions taken by Bayesian DMs.

Note that when the share of boundedly rational DMs vanishes, the bias does not. Indeed, $b_1 > 0$ except for the degenerate case in which $\phi(a^\star(s))$ is constant for all s. The reason for this is that the selectivity of the DM's data does not vanish when agents become Bayesians: they still take higher actions given higher estimates. Thus, decision-making based on data generated by optimal actions results in a bias, albeit a smaller one.

4 Applications

In this section, we apply our results to two settings: project selection and second-price IPV auctions. In both applications the equilibrium is unique, which allows us to derive stronger comparative statics. The main message in both applications is that our naive calibration procedure leads to conservative behavior relative to the behavior of a Bayesian DM. This pessimism takes the form of underbidding in auctions and rejection of marginally good projects.

4.1 Project Selection

The DM is an entrepreneur who selects which projects to implement based on their revenue estimates. We denote the revenue by θ and its estimate by s. Implementing a project entails a cost c. The entrepreneur decides whether to implement the project ($a = 1$) or not ($a = 0$). Thus, the entrepreneur's payoff is $\pi(a, \theta) = a(\theta - c)$. Clearly, she wishes to implement the project if and only if $\theta \geq c$, and so $a^\star(\theta) = 1$ if $\theta \geq c$ and $a^\star(\theta) = 0$ otherwise. Since she believes that the project's revenue is $s - b$ she will launch projects whose estimates are greater than $c + b$.

The entrepreneur bases her decisions on a data set that includes estimates and actual revenues of projects selected by a share α of Bayesian entrepreneurs and a share $1 - \alpha$ of naive entrepreneurs, where $0 \leq \alpha \leq 1$. The data set contains feedback only about implemented projects, namely,[8] $\phi(a) = a$. Denote the entrepreneur's bias by b_α and her cutoff value by s_α such that $s_\alpha = c + b_\alpha$.

[8]Formally, we study the limit case in which $\epsilon \to 0$.

As a benchmark, note that a Bayesian entrepreneur will implement a project if and only if the 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 $s \geq s_B$. The next result compares s_B and s_α under a mild technical assumption and shows that the naive DM is more conservative than the Bayesian DM in the sense that she selects a narrower set of projects, i.e., $s_\alpha \geq s_B$.

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

Claim 1 shows that despite the fact that for sufficiently high estimates the naive entrepreneur's beliefs about the quality of the project are more optimistic than the Bayesian entrepreneur's beliefs (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 $E[s-\theta|s_B] = s_B - c$ as the Bayesian entrepreneur must be indifferent between accepting and rejecting projects whose estimates are equal to her acceptance cutoffs. Were the naive entrepreneur to use this cutoff, her bias would be greater than $s_B - c$ as, since $E[s-\theta|s]$ is increasing, $E[s-\theta|s > s_B] > E[s-\theta|s = s_B]$. Thus, from the naive entrepreneur's perspective, the Bayesian entrepreneur's cutoff is too low.

Jehiel (2018) develops an equilibrium model of project selection in which the baseline setting is similar to the present section's setting: an entrepreneur decides whether to launch a project or not based on a noisy signal. However, his model yields the opposite results: in equilibrium, the agents use an acceptance cutoff which is too low. It is instructive to discuss the different modeling assumptions that lead to this difference.

In Jehiel's model, the entrepreneur has access to the set of all projects that were implemented in the past and their realized revenues. She obtains a signal (interpreted as an impression) about each of these projects, and calculates the expected realized payoff conditional on each signal. Jehiel's entrepreneur believes that, in expectation, the current project's yield will be identical to the expected yield 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 the overoptimism is that the data set 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 payoff of these projects, conditioned on the signal the entrepreneur receives, is higher than the payoff of non-implemented ones. The entrepreneur ignores this selection. Since these projects are of higher quality than the pool of all projects—using 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 data set that the entrepreneur uses to evaluate new projects is selected. However, the different selections yield opposite effects. While Jehiel's entrepreneur uses a data set that contains relatively high states conditional on the signals, our entrepreneur uses a data set that contains high signals conditional on the states. Thus, while Jehiel's entrepreneur concludes that the signals are downward biased, our entrepreneur concludes that they are upward biased.

Alternative interpretation of this application: medical treatment, recommendation systems, and credit provision

While we use the project selection terminology, the analysis in this section is relevant in other contexts as well. We can interpret the DM as a physician who decides which patients to treat based on the results of a medical test and a data set that includes test results and treatment results for previous treated patients. Alternatively, the DM can be viewed as 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 can be interpreted as a credit officer who approves credit applications based on a credit score that is calibrated based on a data set on the return rate of previous successful applications but not of unsuccessful ones.

4.2 Second-price IPV auctions

While our baseline model considers a single DM, its framework can be naturally extended to strategic situations in which multiple players interact. This requires extending the payoff and feedback functions, and making assumptions on how agents' reason about other players' behavior. When a game is dominance solvable the latter become moot.

We now consider a second price IPV auction. There are n 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 drawn from f independently for each bidder. Recall that in a second-price IPV auction bidding one's value is a dominant strategy, i.e., $a^\star(\theta_i) = \theta_i$. We interpret $s_i - b_i$ as agent i's perceived value and, therefore, assume that each bidder i uses the bidding strategy $bid(s_i) = s_i - b_i$, i.e., $a(s_i) = s_i - b_i$. To avoid the problem of potentially negative bids, we assume that $supp(F) = [1, 2]^2$.

We shall assume that a bidder learns the true value of the object if and only if she wins the object. Thus, the higher the bid, the more likely she is to obtain feedback. 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. This means that a bidder obtains feedback if her estimate is the highest, i.e., $\phi(a(s_i, s_{-i})) = (F(s_i))^{n-1}$.

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

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

The next claim shows that the symmetric equilibrium is unique and provides comparative statics with respect to the number of bidders.

Claim 2 *There exists a unique symmetric equilibrium. In this equilibrium, $bid(s) < s$. Moreover, the bias $s - bid(s)$ is increasing in n. Furthermore, $bid(s)$ is decreasing in n.*

The equilibrium uniqueness follows from the fact that the feedback function $\phi(bid(s)) = F(s)^{n-1}$ does not depend on the actual bids in equilibrium, which results in a unique bias that is consistent with our calibration procedure

In a symmetric equilibrium, the bias is strictly positive. This follows from the continuity of the bidding function and the feedback function. This continuity guarantees that non-homogeneous feedback, namely, that different state realizations are observed with different frequencies. Thus, there cannot be a corner 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, note that the feedback function when there are n bidders, $F(s)^{n-1}$, is dominated in the likelihood ratio sense by the feedback function $F(s)^{m-1}$ for $m > n$ bidders.

We now consider the auctioneer's perspective and study her expected revenue. In particular, we compare the case where agents are Bayesian to the case where they use our heuristic. This comparison is not obvious ex ante as our agents' 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. Nonetheless, the next claim shows that, for any number of bidders, the auctioneer's expected revenue is lower when bidders use the heuristic calibration than when bidders are Bayesian.

Claim 3 *The auctioneer's revenue when agents are naive is lower than her revenue when agents are Bayesian.*

When bidders are Bayesian, the winner pays $E[\theta|x_2]$ which is the expected value of the object conditional on the second highest estimate, x_2. Let b_B denote the "Bayesian bias," namely, the expected difference between the second highest estimate and the conditional state realization $b_B = E[x_2 - \theta|x_2]$. Thus, a Bayesian winner pays, on average, $E[x_2] - b_B$. When bidders are naive, they pay $x_2 - b$, where b is the difference between the highest estimate x_1 and $E[x_1 - \theta|x_1]$. Thus, a naive winner pays, on average, $E[x_2] - b$. Since $s - E[\theta|s]$ is increasing in s, and the distribution of the highest signal first-order stochastically dominates the distribution of the second highest signal, it holds that $b > b_B$, and so the naive bidder bids lower than the Bayesian bidder.

Regardless of whether bidders are naive or Bayesian, the bidder with the highest estimate wins the object. Thus, the equilibrium outcome is ex ante efficient. However, the naive calibration procedure affects how the total surplus is divided. On average, naive calibrators are more conservative than Bayesian bidders and, therefore, the auctioneer's share of the surplus is smaller when bidders are naive.

References

[1] Barron, K., Huck, S., and Jehiel, P. (2021): "Everyday econometricians: Selection neglect and overoptimism when learning from others," mimeo.

[2] Chen, M., and Dyer, J. (2007): "Inevitable disappointment in projects selected based on forecasts," In *Hydrocarbon Economics and Evaluation Symposium*. OnePetro.

[3] Esponda, I. (2008): "Behavioral Equilibrium," *American Economic Review,* 98, 1269–91.

[4] Esponda, I., and Pouzo, D. (2016): "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," *Econometrica,* 84, 1093–1130.

[5] Esponda, I., and Vespa, E. (2018): "Endogenous Sample Selection: A Laboratory Study," *Quantative Economics,* 9, 183–216.

[6] Eyster, E. and Rabin, M. (2005): "Cursed Equilibrium," *Econometrica*, 73, 1623–1672.

[7] Jehiel, P. (2018): "Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism," *American Economic Review*, 108, 1582–1597.

[8] Jehiel, P. (2020): "Analogy-based expectation equilibrium and related concepts: Theory, applications, and beyond," mimeo.

[9] Piccione, M. and Rubinstein, A. (2003): "Modeling the Economic Interaction of Agents with Diverse Abilities to Recognize Equilibrium Patterns," *Journal of the European Economic Association,* 1, 212–223.

[10] Spiegler, R. (2016): "Bayesian Networks and Boundedly Rational Expectations," *Quarterly Journal of Economics*, 131, 1243–1290.

[11] Spiegler, R. (2020): "Behavioral Implications of Causal Misperceptions," *Annual Review of Economics* 12, 81–106

5 Appendix: Proofs

Proof of Proposition 1. Since $E(\theta) = E(s)$, we have that

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

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

$$T(b) := \frac{\int_{-\infty}^{\infty} f(s)(\phi(a^\star(s-b)))[s-E(\theta|s)]ds}{\int_{-\infty}^{\infty} f(s)\phi(a^\star(s-b))ds + \delta}. \tag{2}$$

It is easy to verify that $T(b)$ is well-defined and continuous. To establish equilibrium existence, we need to find a bias $b^\star$ such that $T(b^\star) = b^\star$.

Since $lim_{b\to\infty}\phi(a^\star(s-b)) = k_l \in [0,1]$ for every s, we have that

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

Since $lim_{b\to-\infty}\phi(a^\star(s-b)) = k_h \in [0,1]$ for every s, we have that

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

Since the denominator of (2) is strictly positive in the limits, it follows that $lim_{b\to\infty}T(b) = lim_{b\to-\infty}T(b) = 0$. Since $T(b)$ is continuous, there exists $b^\star$ for which $T(b^\star) = b^\star$.

Proof of Proposition 2. Fix an arbitrary b. Consider the RHS of (2). Let $s^\star$ denote an arbitrary estimate that satisfies $s^\star - E[\theta|s^\star] = 0$. Observe that $s - E[\theta|s]$ is non-positive for $s < s^\star$ and non-negative for $s > s^\star$. Since $a^\star$ is increasing in s, and ϕ is increasing in a we obtain that

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

and

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

The sum of the RHS of the two inequalities is 0 as $E(\theta) = E(s)$. Thus, the numerator of the RHS of (2) (which is equal the the sum of the LHSs of the two inequalities) is positive. Clearly, its denominator is strictly non-negative. Hence, $T(b) \geq 0$ for any b. This completes the proof as, in equilibrium, $T(b) = b$.

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

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

Let

$$\alpha = \frac{\int_{-\infty}^{\infty} f(s)\tilde{\phi}(a^\star(s-b))ds}{\int_{-\infty}^{\infty} f(s)\phi(a^\star(s-b))ds}$$

and note that the above inequality holds if and only if

$$(3) \qquad \int_{-\infty}^{\infty} f(s)(\alpha\phi(a^\star(s-b)) - \tilde{\phi}(a^\star(s-b)))[s-E(\theta|s)]ds \geq 0.$$

Since $\int_{-\infty}^{\infty} f(s)[s-E(\theta|s)]ds = 0$ and $s-E(\theta|s)$ is increasing in s, (3) holds if $\alpha\phi(a^\star(s-b)) - \tilde{\phi}(a^\star(s-b))$ is increasing in s. The latter holds if and only if $\alpha \geq \frac{\tilde{\phi}'(a^\star(s-b))}{\phi'(a^\star(s-b))}$. Since ϕ dominates $\tilde{\phi}$ in the likelihood ratio sense, $(\frac{\tilde{\phi}(a^\star(s-b))}{\phi(a^\star(s-b))})' \leq 0$. Since the derivative of the feedback function w.r.t., s is positive, this implies that $\frac{\tilde{\phi}'(a^\star(s-b))}{\phi'(a^\star(s-b))} \leq \frac{\tilde{\phi}(a^\star(s-b))}{\phi(a^\star(s-b))}$ and, as a result, $\alpha \geq \frac{\tilde{\phi}'(a^\star(s-b))}{\phi'(a^\star(s-b))}$.

Proof of Proposition 4. A necessary condition for an equilibrium is $T_\alpha(b_\alpha) = b_\alpha$, where, in a similar manner to (2), $T_\alpha(b_\alpha)$ can be written as

$$(4) \quad T_\alpha(b_\alpha) := \frac{\int_{-\infty}^{\infty} f(s)[(1-\alpha)\phi(a^\star(s-b_\alpha)) + \alpha\phi(a^\star(s-b_B(s))]E[s-\theta|s]ds}{\int_{-\infty}^{\infty} f(s)[(1-\alpha)\phi(a^\star(s-b_\alpha)) + \alpha\phi(a^\star(s-b_B(s))]ds + \delta}.$$

We will show that (4) is decreasing in α when $T_\alpha(b_\alpha) = b_\alpha$. Thus, for the equilibrium condition to hold, b_α must become lower when α becomes higher. The derivative of (4) with respect to α is negative if and only if

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

We can rewrite this inequality as

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

Since $b_B(s) = E[s-\theta|s]$ is increasing in s and $b_\alpha = T_\alpha(b_\alpha)$ is a weighted average of

$b_B(s)$, we obtain that either $\phi(a^\star(s - b_B(s)) - \phi(a^\star(s - b_\alpha)) \leq 0$ or $b_B(s) - b_\alpha \leq 0$, which implies the above inequality. We can conclude that when α is increased, the RHS is smaller than the LHS for b_α that satisfies $T_\alpha(b_\alpha) = b_\alpha$. Since $T_\alpha(b_\alpha)$ is continuous in b_α, we obtain that the bias b'_α for which $T_\alpha(b'_\alpha) = b'_\alpha$ is lower than b_α when α increases.

Proof of Claim 1. First, note that, in equilibrium, it must be that

$$s_\alpha - b_\alpha = c. \tag{5}$$

Since the entrepreneur implements all projects whose estimate is greater than s_α and $\phi(a) = a$, it follows that $b_\alpha = E[s|s \geq s_\alpha] - E[\theta|s \geq s_\alpha]$. Hence, the equilibrium cutoff s_α satisfies

$$s_\alpha - E[s|s \geq s_\alpha] = c - E[\theta|s \geq s_\alpha]. \tag{6}$$

Since $E[\theta|s]$ is increasing in s, the RHS of 6is 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 $b_\alpha \geq 0$. Recall that $b_\alpha = E[s - \theta|s \geq s_\alpha]$. Since $E(s) = E(\theta)$ and $s - E[\theta|s]$ is increasing is s it holds that $E[s - \theta|s \geq s_\alpha] \geq 0$ 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 (6) becomes

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

Hence, for (6) 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 (5).

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

$$b = \frac{\int_{\underline{s}}^{\overline{s}} f(s)(F(s)^{n-1})[s - E(\theta|s)]ds}{\int_{\underline{s}}^{\overline{s}} f(s)(F(s)^{n-1})ds + \delta}. \tag{7}$$

As the RHS of (7) is independent of b, it has a solution which is unique.

By Proposition 1, $b \geq 0$ and, therefore, $bid(s) \leq s$.

When the number of bidders is n, the feedback function is $\phi_n(a(s-b_n)) = F(s)^{n-1}$. Thus, $\frac{\phi_{n+1}(a(s-b_{n+1}))}{\phi_n(a^\star(s-b_n))} = F(s)$ is increasing in s and, by Proposition 2, the proof is complete.

Proof of Claim 3. Denote the k-th order statistic of a sample of n independent draws from $f(s)$ by x_k and denote its density function by f_k. The expected value of the object, conditioned on winning the auction, is $\int_{\underline{s}}^{\overline{s}} E[\theta|s]f_n(s)ds$. The expected value of the estimate, conditioned on winning the auction, is the expected value of the n'th order statistic, $\int_{\underline{s}}^{\overline{s}} sf_n(s)ds$. Thus, the equilibrium bias is

$$b = \int_{\underline{s}}^{\overline{s}} sf_n(s)ds - \int_{\underline{s}}^{\overline{s}} E[\theta|s]f_n(s)ds. \tag{8}$$

Since players subtract the bias from their estimates when bidding, and the winner pays the second highest bid, the auctioneer's expected revenue is the $n-1$-th order statistic net of the bias,

$$\int_{\underline{s}}^{\overline{s}} sf_{n-1}(s)ds - b. \tag{9}$$

A Bayesian bidder would bid the expected value of the object given the signal she receives. Thus, if all bidder were Bayesian, the auctioneer's revenue would be the the expected value of θ conditioned on receiving the second highest estimate,

$$\int_{\underline{s}}^{\overline{s}} E[\theta|s]f_{n-1}(s)ds \tag{10}$$

By (8), (9) and (10), the auctioneer's revenue is higher when agents are Bayesian if and only if

$$\int_{\underline{s}}^{\overline{s}} (E[\theta|s] - s) f_{n-1}(s) ds > \int_{\underline{s}}^{\overline{s}} (E[\theta|s] - s) f_n(s) ds. \tag{11}$$

Condition (11) holds since x_n first-order stochastically dominates x_{n-1}, and $E[\theta|s] - s$ is decreasing in s by assumption.