Mental Models and Endogenous Learning

Hakan Özyılmaz[*]
(Toulouse School of Economics)

February 24, 2023

Abstract

This paper experimentally studies how people learn about their environment when their subjective understanding of the environment, their mental model, is misspecified. We use people's tendency to hold optimistic beliefs about their abilities to generate model misspecification and investigate the implications of overconfidence as a source of misspecified mental model on learning about own ability and a fundamental. Consistent with the theoretical predictions, overconfident subjects develop pessimistic beliefs about the fundamental and take growingly suboptimal actions. Inconsistent with the theoretical prediction, endogenous feedback does not exacerbate the extent of suboptimal behavior. Investigating how subjects learn about their own ability reveals that abundant feedback "weakens" misspecified mental models. The "weakening" of mental models is more pronounced with endogenous feedback and explains why endogenous feedback may not exacerbate the extent of suboptimal behavior.

[*]I thank Erik Eyster, Emanuel Vespa, Gary Charness, Ignacio Esponda, Paul Heidhues, Maria Kogelnik, Ryan Oprea, Philipp Strack, Sevgi Yuksel, and seminar participants at ESA North America 2022 for helpful comments. Research support from UCSB Department of Economics and Toulouse School of Economics is gratefully acknowledged. Özyılmaz: Toulouse of Economics, 1 Esplanade de l'Université, T.622, 31080 Toulouse, France, hakan.ozyilmaz@tse-fr.eu.

1 Introduction

People regularly construct mental models of their environment to guide their reasoning, to make inferences and to understand how their actions map into outcomes. Accumulating evidence from psychology and economics documents that people frequently create mental models that fall short of being an accurate representation of their decision environment, as they struggle gathering, attending, and processing crucial information in their environment (Hanna, Mullainathan and Schwartzstein (2014), Handel and Schwartzstein (2018)). An important question is then if the constructed models do not admit the possibility of truth, what are the implications of this "misspecification" for how people learn about their environment, and subsequently make decisions that are informed by their learning? In particular, do more data necessarily mean that people will correctly learn a fundamental variable and take the first-best action?

We set out to answer this question using a carefully designed laboratory experiment. In the experiment, subjects repeatedly make investment decisions on a fixed project over 1000 periods, and in each period observe a noisy signal on the profit they generate. Subjects do not know the quality of their assigned project, but correctly know that the assignment is random. Moreover, there are complementarities between the investment amount and the expected project quality – projects with higher expected quality require higher investment at the optimum. Another key determinant of the profit is an ability parameter for the subject that positively affects the profit. In particular, we assign each subject an ability score based on their ranking on an "IQ test" that they take at the beginning of the experiment. Tying an ego-relevant ability parameter to the profit and not resolving the uncertainty over the ability parameter provides scope for model misspecification in how subjects perceive the profit equation. In our controlled environment, subjects frequently create misspecified models of the profit equation because of their overestimation of their ranking on the IQ test. We find that 34% of our subjects assign 100% likelihood to ability scores that are *strictly* above their true abilities.

The theory makes sharp predictions about how subjects with misspecified models of the profit equation should take actions throughout the experiment. Heidhues, Kőszegi and Strack (2018) show that if agents with misspecified models do not update their prior on their ability and learn from feedback in a Bayesian fashion, they should on average make growingly suboptimal investments in this environment. Intuitively, this is because agents regularly experience less-than-expected profit because of their overconfidence and explain the less-than-expected output by developing pessimistic beliefs about the project quality. Consistent with this prediction, we find that overconfident subjects with misspecified models of the profit equation make growingly suboptimal investments throughout the experiment's time horizon. By the last period of the experiment, overconfident subjects significantly under-invest relative to the first-best in their projects compared to the subjects who have a correctly specified model of the profit equation.

A further theoretical insight in Heidhues, Kőszegi and Strack (2018) that directly applies to our setting is that endogenous learning exacerbates the extent of suboptimal behavior when agents have misspecified models. This is because suboptimal behavior generated through model misspecification further depresses profit. Agents rationalize the additionally depressed profits by even more pessimistic beliefs about the project quality. These more pessimistic beliefs about the project quality then lead overconfident subjects to curb investment even further because of the complementarities in project qualities and investment amounts. In order to test this prediction we create two treatments **Exogenous** and **Endogenous** where we manipulate the endogeneity of feedback. While subjects in **Endogenous** see their investment decisions immediately implemented in each period and receive feedback that comes from the "profit distribution" they select in that period, subjects in **Exogenous** do not see their investment decisions immediately implemented and instead receive feedback from a fixed pre-announced "profit distribution" throughout the experiment. Although overconfident subjects in both treatments make growingly suboptimal investments, we do not find that endogenous feedback exacerbates the suboptimal investment behavior. Although subjects in **Exogenous** take actions that are statistically indistinguishable from the action a myopically optimizing Bayesian agent would take on average by the last period of the experiment, subjects in **Endogenous** sharply deviate from this Bayesian benchmark.

Investigating how subjects in **Exogenous** and **Endogenous** learn about their abilities reveals insights into the deviation from the theoretical prediction and the Bayesian benchmark. Although Bayesian learners would learn virtually nothing about their abilities in our experiment, a comparison of elicited prior and posterior means clearly indicate that overconfident subjects have become less overconfident by the end of the experiment. This "weakening" of mental models, including the truth within the set of possibilities, in the face of abundant objective feedback is consistent with previous work (Esponda, Vespa and Yuksel (2020)). Interestingly, we find that this "weakening" is more pronounced for subjects who face endogenous feedback although this effect is not significant at conventional levels.[1]

Substantial evidence in psychology suggests that on average people have unrealistically positive views of their traits (e.g. Weinstein (1980), Svenson (1981), Moore and Healy (2008)). A large literature in economics investigates how such overconfident beliefs about own traits and prospects lead people to make suboptimal decisions (excess entry decisions in the laboratory, Camerer and Lovallo (1999); over-trading by retail investors, Barber and Odean (2001); over-investment by CEOs, Malmendier and Tate (2005), Malmendier and Tate (2008)). As the evidence on the material costs of overconfidence accumulates, theoretical and laboratory studies started exploring how people produce and maintain overconfident beliefs about their abilities. While the theoretical literature documented how overconfidence may arise as a result of biased memory, ego utility , motivational and signalling values (Bénabou and Tirole (2002), Köszegi (2006)), laboratory studies confirmed these mechanisms (Zimmermann (2020), Möbius et al. (2014), Chen and Schildberg-Hörisch (2019)) and further documented how biased processing of objective noisy feedback prevent people from learning their true abilities (Eil and Rao (2011), Ertac (2011), Coutts (2019)).

Our investigation pushes the literature on overconfident agents' learning processes forward in three fundamental ways. First, while the previous literature focuses on settings where there is a single source of uncertainty (e.g. the agent's ability) that generates noisy feedback, we focus on a setting where the noisy feedback features two sources of uncertainty (e.g. the agent's ability and an external stable fundamental). Many economically important settings feature such multidimensional sources of uncertainty and a lower dimensional feedback. Examples include an employee not knowing her marginal return to effort, deciding how hard she wants to work and observing the output of her efforts; a team member not knowing the ability of his teammate, deciding how much of the work to delegate and observing the joint output produced by his team. Second, while the focus of the previous literature is on how overconfident people learn about their abilities through noisy feedback, we focus on how overconfident people learn not about their ability, but about an external decision-relevant variable. There is ample evidence documenting that people are good at supplying overconfident beliefs about their own abilities, however, in the environments that we are investigating there is little evidence if the supply of such beliefs biases

[1]However, the greater reduction in overconfidence is consistent with our main finding that endogenous learning does not exacerbate suboptimal investment

the way people learn about their environment. Third, while the previous literature is interested in exogenous learning situations where individuals learn about their abilities without choosing actions, our focus is on endogenous learning situations where individuals are provided opportunities to take actions that allow them to sample feedback from distinct distributions. It is clear that most learning environments have this “experimentation" feature (e.g. an employee might try working very hard or very little to get a more precise signal of her own ability) and hence, arguably, is of greater economic relevance than exogenous learning situations.

2 Theoretical Framework

In this section, we present our theoretical framework and discuss its main predictions. The model has two main goals. First, it illustrates an environment where overconfident agents who have a misspecified model of their environment underestimate an external fundamental while such underestimation does not necessarily arise for agents who have a correctly specified model. Second, it highlights how endogenous learning might exacerbate the extent of underestimation for overconfident agents.

2.1 Overview

We focus on a simple decision environment where the agent is uncertain about her ability and an external fundamental that affects the output she generates. The agent periodically takes an action and receives a noisy feedback on the output she generates. More specifically, let $a \in A = \{20, 40, 60, 80, 100\}$ represent the agent's unchanging ability and $\phi \in \Phi = [0, 100]$ represent the unobservable unchanging fundamental. We assume the fundamental is randomly drawn from the uniform distribution $\pi_0 : \Phi \longrightarrow \mathbb{R}_+$ before the agent starts making her decisions and is independent of the agent's ability a. The agent chooses an action $e_t \in E = [0, 100]$ in each period and produces an output $y(e_t, a, \phi)$ with partial derivatives $y_a \geq 0, y_\phi > 0$. In particular, we assume that the output has a simple functional form

$$y(e_t, a, \phi) = (a + e_t)\phi - \frac{e_t^2}{2}$$

After each action, the agent observes a noisy feedback f_t on the output she generates. The feedback is distributed Bernoulli with mean $\mu(e_t, a, \phi)$ that corresponds to the normalization of the output function:

$$\mu(e_t, a, \phi) = \frac{y(e_t, a, \phi) - \underline{y}}{\bar{y} - \underline{y}}$$

where $\bar{y} = \max_{e,a,\phi} y(e, a, \phi)$ and $\underline{y} = \min_{e,a,\phi} y(e, a, \phi)$.

2.2 Objective Model

Fix an ability level for the agent, a_o, and her teammate, ϕ_o. For each effort level e_t, there is an objective feedback distribution $Q_o(\cdot|e_t)$ that is a Bernoulli density with mean $\mu(e_t, a_o, \phi_o)$.

2.3 Mental Model

The mental model represents the set of feedback distributions the agent considers possible a priori. For a fixed objective decision problem Q_o, a mental model is a tuple

$$\mathcal{Q} = \langle \Theta, (Q_\theta)_{\theta \in \Theta} \rangle$$

where $\Theta \subset A \times \Phi$ is the agent's parameter set and $Q_\theta(\cdot|e_t)$ is the action-dependent feedback distributions parametrized by $\theta = (a, \phi) \in \Theta$. While the action-dependent objective feedback distribution $Q_o(\cdot|e_t)$ represents the true environment, the mental model represents the agent's perception of their environment.

We assume that the agent correctly believes that the map from actions to probability distributions over feedback is fixed and depends only on their current action, but they are uncertain about the distribution each action induces. The agent's uncertainty about what the true environment $Q_o(\cdot|e_t)$ is captured by their mental model $\langle \Theta, (q_\theta)_{\theta \in \Theta} \rangle$ and a joint density function $\Pi_0 : A \times \Phi \longrightarrow \mathbb{R}_+$ that describes the agent's prior belief. Following the previous literature, we call the agent's mental model correctly specified if the true parameter vector lies in the support of the agent's prior beliefs ($\theta_o = (a_o, \phi_o) \in \Theta$), and otherwise call it misspecified.

We assume that the agent correctly believes that the fundamental is independently drawn from their own ability and from the uniform distribution π_0 with the support Φ. Due to independence, we can decompose the agent's prior density $\Pi_0(a, \phi) = p_0(a)\pi_0(\phi)$ where p_0 is a probability mass function that describes agent's prior belief about own ability. In this environment, we assume that the agent chooses myopically optimal actions at each period, aiming to maximize the probability of getting a "positive" feedback and learns from feedback using Bayes' rule.

2.4 Overconfidence as a Misspecified Mental Model

Since the agents in this framework have potentially misspecified models of their environment, we apply the solution concept proposed in Esponda and Pouzo (2016) to derive the set of possible limit points of the agent's learning process.

The equilibrium requires agents' beliefs to put probability 1 on the set of subjective feedback distributions that are “closest" to the objective distribution. Building on Berk (1966), Esponda and Pouzo (2016) shows that the correct notion of “distance" is the Kullback-Leibler divergence in statistics. It represents a “distance" between the objective output distribution $Q_o(\cdot|e)$ and the family of parametrized subjective distributions $(Q_\theta(\cdot|e))_{\theta\in\Theta}$ for a fixed action e:

$$K(e, \theta) = E_{Q_o(\cdot|e)} log\left[\frac{Q_o(f|e)}{Q_\theta(f|e)}\right]$$

Both objective and subjective mental models belonging to the family of Bernoulli distributions, the KL divergence in our context is simply

$$\begin{aligned} K(e, \theta) &= E_{Q_o(\cdot|e)}\left[f\, log\frac{\mu(e, a_o, \phi_o)}{\mu(e, a, \phi)} + (1-f) log\frac{1-\mu(e, a_o, \phi_o)}{1-\mu(e, a, \phi)}\right] \\ &= \mu(e, a_o, \phi_o) log\frac{\mu(e, a_o, \phi_o)}{\mu(e, a, \phi)} + (1-\mu(e, a_o, \phi_o)) log\frac{1-\mu(e, a_o, \phi_o)}{1-\mu(e, a, \phi)} \end{aligned}$$

The set of closest parameter values for the agent given an effort decision e can then be described as

$$\hat{\Theta}(e) = \arg\min_{\theta\in\Theta} K(e, \theta)$$

The interpretation is that $\hat{\Theta}(e) \subset \Theta$ is the set of parameter values that the agent can believe to be possible after observing feedback consistent with the effort decision e.

A pure strategy Berk-Nash equilibrium of a single agent problem is then a pair of action and a belief (e^*, Π^*) that satisfies

i) $e^* \in \arg\max_{e \in E} E_{\bar{Q}_\Pi(\cdot|e)} f$ where $\bar{Q}_\Pi = \int_\Phi Q_\theta \Pi^*(d\theta)$

ii) $\Pi^* \in \Delta(\hat{\Theta}(e^*))$

Defining Overconfidence. In this environment, we define an overconfident agent to be one whose prior on their own ability assigns zero mass on their true ability and is supported by abilities that are greater than their own ability i.e. $a > a_o$ for any $a \in \text{supp} p_0(a)$. Hence an overconfident agent's mental models are misspecified as $\theta_o = (a_0, \phi_o) \notin \Theta$.

2.4.1 Exogenous Learning

We first look at how an overconfident agent learns the fundamental when their action is fixed and when they are provided with infinite feedback. Fix an effort decision $\bar{e}$. Assume that for each $a \in \text{supp} p_0(a)$, there is $\phi_a \in \text{supp} \pi_0(\phi)$ such that $\mu(\bar{e}, a_o, \phi_o) = \mu(\bar{e}, a, \phi_a)$. This assumption ensures that for any fixed action, the agent can always find a fundamental that explains the observed distribution of feedback irrespective of what they believe their own ability level to be. This implies KL divergence is minimized at 0 for all (a, ϕ_a) where $a \in \text{supp} p_0(a)$ and generates the following set of KL minimizers

$$\hat{\Theta}(\bar{e}) = \{(a, \phi_a = \frac{a_o + \bar{e}}{a + \bar{e}} \phi_o) | a \in \text{supp} p_0(a)\}$$

Lemma 1. *Suppose that the agent takes a fixed action $\bar{e}$ in all periods and for each $a \in \text{supp} p_0(a)$, there exists $\phi_a \in \text{supp} \pi_0(\phi)$ such that*

$$\mu(\bar{e}, a_o, \phi_o) = \mu(\bar{e}, a, \phi_a)$$

Then the agent's beliefs on (a, ϕ) almost surely converges and concentrates on $\hat{\Theta}(\bar{e})$.

Proof. See Berk (1966). ∎

Proposition 1. *Let $\Pi_\infty^{\bar{e}}$ be the limiting posterior distribution on (a, ϕ) when the agent repeatedly chooses effort $\bar{e}$. Then $\Pi_\infty(a, \phi) = p_0(a)\mathbb{1}_{\hat{\Theta}(\bar{e})}(a, \phi)$.*

Proof. By Lemma 1, $\Pi_\infty^{\bar{e}}(a, \phi) = 0$ whenever $(a, \phi) \notin \hat{\Theta}(\bar{e})$. Moreover, $\hat{\Theta}(\bar{e})$ is finite since $p_o(a)$ has finite support. Take (a_1, ϕ_{a_1}) and $(a_2, \phi_{a_2}) \in \hat{\Theta}(\bar{e})$,

$$\begin{aligned}
\lim_{t\to\infty} \frac{\Pi_t(a_1, \phi_{a_1} \mid f_1, .., f_t)}{\Pi_t(a_2, \phi_{a_2} \mid f_1, .., f_t)} &= \lim_{t\to\infty} \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \frac{Q_{\theta_1}(f_1, .., f_t \mid \bar{e})}{Q_{\theta_2}(f_1, .., f_t \mid \bar{e})} \\
&= \lim_{t\to\infty} \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \frac{\mu(\bar{e}, a_1, \phi_{a_1})^{tf}(1 - \mu(\bar{e}, a_1, \phi_{a_1}))^{t(1-f)}}{\mu(\bar{e}, a_2, \phi_{a_2})^{tf}(1 - \mu(\bar{e}, a_2, \phi_{a_2}))^{t(1-f)}} \\
&= \lim_{t\to\infty} \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \frac{\mu(\bar{e}, a_o, \phi_o)^{tf}(1 - \mu(\bar{e}, a_o, \phi_o))^{t(1-f)}}{\mu(\bar{e}, a_o, \phi_o)^{tf}(1 - \mu(\bar{e}, a_o, \phi_o))^{t(1-f)}} \\
&= \lim_{t\to\infty} \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \\
&= \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \\
&= \frac{p_0(a_1)\pi_0(\phi_{a_1})}{p_0(a_2)\pi_0(\phi_{a_2})} \\
&= \frac{p_0(a_1)}{p_0(a_2)}
\end{aligned}$$

Hence $\Pi_\infty^{\bar{e}}(a, \phi_a) = p_0(a)$ for each $(a, \phi_a) \in \hat{\Theta}(\bar{e})$ ∎

Corollary 1. *A overconfident agent's learning process leads him to underestimate the fundamental i.e. whenever min* $\operatorname{supp} p_0(a) > a_o$, $E_{\Pi_\infty^{\bar{e}}}[\phi] < \phi_o$.

Proof. For any $a \in \operatorname{supp} p_0(a)$, $\phi_a = \frac{a_o+\bar{e}}{a+\bar{e}}\phi_o < \phi_o$. Hence $E_{\Pi_\infty^{\bar{e}}}[\phi] < \phi_o$. ∎

2.4.2 Endogenous Learning

We now look at what the overconfident agent comes to believe about the fundamental when he is allowed to change his ction in each period in response to his beliefs. In particular, we are interested in if his inferences about the fundamental improve when he chooses myopically optimal actions in each period. A Berk-Nash equilibrium for an overconfident agent with prior

$p_0(a)$ is a pair $(\Pi_\infty(a, \phi), e^*)$ such that

$$i) \;\; e^* = E_{\Pi_\infty}[\phi] \tag{1}$$

$$ii) \;\; \Pi_\infty(a, \phi) = p_0(a)\mathbb{1}_{\hat{\Theta}^J(e^*)}(a, \phi) \tag{2}$$

where

$$\hat{\Theta}(e^*) = \{(a, \phi_a = \frac{a_o + e^*}{a + e^*}\phi_o)|a \in \text{supp}p_0(a)\} \tag{3}$$

Lemma 2. *Assume the equations (1) and (2) have a solution. Then $e^* = E_{\Pi_\infty}[\phi] < \phi_o$.*

Proof. Note that $\frac{a_o+e^*}{a+e^*} < 1$ for any $a \in \text{supp}p_0(a)$. Then

$$e^* = E_{\Pi_\infty}[\phi] = \sum_{a \in \text{supp}p_0(a)} p_0(a)\frac{a_o + e^*}{a + e^*}\phi_o < \phi_o \tag{4}$$

■

The above lemma shows that the equilibrium action is lower than the fundamental.

Proposition 2. *Let $\Pi_\infty^{\bar{e}}$ be the limiting posterior distribution when the agent exogenously learns from data with fixed action $\bar{e}$ and let Π_∞ be the limiting posterior distribution when the agent learns endogenously with the optimal action. Then $E_{\Pi_\infty}[\phi] < E_{\Pi_\infty^{\bar{e}}}[\phi]$ whenever $\bar{e} \geq \phi_o$.*

Proof. Note that for $\phi_a(\bar{e}) = \frac{a_o+\bar{e}}{a+\bar{e}}\phi_o$

$$sgn(\frac{\partial \phi_a}{\partial \bar{e}}) = sgn(a - a_o) > 0 \tag{5}$$

Hence for an overconfident agent ($min \;\; \text{supp}p_0(a) > a_o$), learning from a higher fixed effort decision leads to a higher belief on ϕ for each KL minimizing pair (a, ϕ_a). Since $e^* < \phi_o$, whenever $\bar{e} \geq \phi_o$

$$\frac{a_o + e^*}{a + e^*}\phi_o < \frac{a_o + \phi_o}{a + \phi_o}\phi_o \leq \frac{a_o + \bar{e}}{a + \bar{e}}\phi_o \quad \forall a \in \text{supp}p_0(a) \tag{6}$$

Therefore,

$$\sum_{a \in \operatorname{supp} p_0(a)} p_0(a) \frac{a_o + e^*}{a + e^*} \phi_o < \sum_{a \in \operatorname{supp} p_0(a)} p_0(a) \frac{a_o + \bar{e}}{a + \bar{e}} \phi_o \tag{7}$$

$$E_{\Pi_\infty}[\phi] < E_{\Pi_\infty^{\bar{e}}}[\phi] \tag{8}$$

■

Proposition 2 implies that if the agent starts off with an effort level at or above the optimal effort level, then the opportunity to change his effort decision in response to his inferences leads to more *incorrect* long-run expectations than if he could not change his effort decision. We call this type of learning *self-defeating*.

Corollary 2. *When $\bar{e} \geq \phi_o$, an overconfident agent's long-run optimal effort decision e^* is more inaccurate when he is allowed to change his effort decision in response to his inferences than a hypothetical optimal effort decision $\bar{e}^*$ he would like to take when he could not change his effort decision.*

Proof. Since the optimal effort level is equal to the expectation of ϕ under the joint feedback, we have $e^* = E_{\Pi_\infty}[\phi] < E_{\Pi_\infty^{\bar{e}}}[\phi] = \bar{e}^* < \phi_o$ ■

Intuitively, the overconfident agent is "surprised" by the negative feedback he observes when he collects sufficient data to identify the feedback distribution that he faces when he repeatedly takes an action greater than the first-best action, that is, the action he would take if he were to know the fundamental. The reason is that the feedback he receives increases in his ability and hence he expects higher feedback than actually realized. Once he identifies a feedback distribution that is lower than his expectation, he attributes the low output he generated to the fundamental being lower than his expectation. The beliefs that the agent develop lead him to exert lower than the first best action as his incentives to take higher actions increase in the fundamental. Since the action he takes increases the probability of receiving positive feedback when he chooses an action that is lower the first-best level, he decreases his probability of receiving positive feedback by choosing a lower action. This provides further negative feedback to the overconfident agent that "surprises" him, he explains these further negative feedback by lowering his expectations about the fundamental even further. This process continues until the overconfident agent is no longer "surprised" about the feedback he receives.

2.5 Hypotheses

Hypothesis 1. When the actions are exogenous and fixed at a level $\bar{e}$ greater than the first-best level and if beliefs converge:

- *an overconfident agent's expectation on the fundamental converges to a point that is less than the fundamental i.e.* $E_{\Pi_\infty^{\bar{e}}}[\phi] < \phi_o$

Our second hypothesis is that learning is *self-defeating* for overconfident agents i.e. when provided an opportunity to revise their actions in response to their inferences, overconfident agents' expectations are further away from the truth.

Hypothesis 2. When the actions are endogenous and if beliefs and actions converge:

- *an overconfident agent's expectation converges to* $E_{\Pi_\infty}[\phi]$ *that satisfies* $E_{\Pi_\infty}[\phi] < E_{\Pi_\infty^{\bar{e}}}[\phi] < \phi_o$
- *an overconfident agent's action converges to a point that is less than the first-best level*

Our third hypothesis involves learning about own ability.

Hypothesis 3. Irrespective of the endogeneity of actions, if beliefs converge:

- *an overconfident agent's expectation on his own ability converges to a point that is identical to his prior expectation on his own ability i.e.* $E_{\Pi_\infty}[a] = E_{\Pi_0}[a]$

2.6 Correctly Specified Mental Models

When the agent's prior about his own ability assigns some mass to his true ability a_o, his mental model is correctly specified. In this instance, the agent's beliefs do not need to converge to his true ability and the true fundamental as he can only exactly identify the feedback distribution he faces while unable to pin down the underlying parameters (a and ϕ) of that distribution. The predictions for such agents are ambiguous and prior-specific. The two examples below show that an agent who has a correctly specified mental model yet who expects his ability to be greater than his actual ability might 1) grow pessimistic or optimistic about the fundamental and 2) endogenous learning might either exacerbate or alleviate the extent of mislearning.

Example 1 (*Almost Misspecified*). Suppose the agent's prior on his own ability $p_0(a)$ is such that $a \geq a_o$ for any $a \in \text{supp} p_0(a)$ with $p_0(a_o) \in (0, 1)$. Clearly, $E_{p_0}[a] > a_o$. So the agent has a correctly specified model where he expects his ability to be strictly better than his actual ability. It is easy to see that Lemma 2 and Proposition 2 are still valid for this agent. Hence the agent is pessimistic about the fundamental when he learns under a fixed action and he exhibits self-defeating learning when he is allowed to change his action in response to his inferences.

Example 2 (*Almost Symmetric Prior Around* a_0). Suppose the agent's prior on his own ability $p_0(a)$ is such that $\text{supp} p_0(a) = \{a_+, a_o, a_-\}$ where $a_+ - a_o = a_o - a_-$ with $p_0(a_+) = \bar{p} + \varepsilon$, $p_0(a_o) = \bar{p}$, $p_0(a_-) = \bar{p} - \varepsilon$. Again, $E_{p_0}[a] > a_o$. Since the KL minimizing belief $\phi_a = \frac{a_o + \bar{e}}{a + \bar{e}} \phi_o$ is strictly decreasing and convex in a, $E_{\Pi^{\bar{e}}_\infty}[\phi] > \phi_o$. Hence the agent becomes optimistic about his teammate's ability when he learns under a fixed effort decision. We can use a little bit more algebra to show that the agent's expectations about his teammate move closer to the truth when he learns endogenously compared to the situation where he learns exogenously at the first-best effort level. Hence his endogenous learning is self-correcting.

3 Experimental Design

The objective of the design is to construct a decision environment in which i) agents are likely to form misspeficied mental models and ii) see if and how these mental models misguide learning about payoff-relevant decision variables.

Overview

The experiment consists of five parts. At the beginning of each part of the experiment, we provide subjects with the instructions, familiarize them with the interface and test their understanding of the rules of the experiment through a series of understanding quizzes. In the first part of the experiment, we measure the "ability" of our subjects using Raven matrices framed as an IQ test. In the second part of the experiment, we elicit subjects' beliefs about their relative performance on the IQ test compared against 19 randomly selected participants who participated in a pilot session.

The third part of the experiment is the main part where each subject is randomly assigned to one of our treatments. At the beginning of this part of the experiment, we assign subjects different fundamentals. More specifically, we frame the decision environment as subjects acting as project managers for a company where projects correspond to fundamentals.[2] Subjects' relative ranking on the IQ test and their assigned fundamentals jointly determine the probability of receiving a positive feedback. The assigned fundamental for each subject remains constant until the end of this part. In each period, subjects are required to submit actions framed as investment recommendations on their assigned projects. In order to help them with their decisions, subjects are provided with special calculators that take their beliefs about their own ability parameter "a" as an input and calculate the myopically optimal actions. In treatment **Exogenous**, subjects actions are *not* implemented to generate feedback but they are implemented for their payment. In treatment **Endogenous**, subjects' actions are implemented to generate feedback and also for the calculation of their payments.

In the fourth part of the experiment, we re-elicit subjects' beliefs about their ranking on the IQ test they have taken at the beginning of the experiment. In the fifth and the final part of the experiment, subjects complete a survey where they are asked to provide basic demographic information. In both treatments, subjects' payoffs are determined by the sum of the amount they made in a randomly selected part (either \$25 or \$0) and a show-up fee of \$10.

3.1 Part 1: Establishing Ability Parameters

The goal of this part of the experiment is to establish an ability parameter for each subject. We measure subjects' ability parameters using Raven's matrices. Subjects are introduced to Raven's matrices as a test of intelligence to accentuate the ego-relevance of the task and to provide scope for overconfidence. We ask each subject to solve the same 10 Raven's matrices, present them in the same order and provide the subjects with 10 minutes to finish the test. Once subjects finish the test, we compare the number of correctly answered questions to the performance of 19 randomly selected subjects who took the exact same IQ test in a pilot session of the experiment. Each subject

[2] Subjects' task is to recommend investment decisions (actions) to the company that is to be invested into their assigned projects (fundamentals). Their goal is to maximize their profit (output) from the project. In each period after they make an investment recommendation, they get an evaluation from the company if their profit for that period beats the company's profit expectations or not (Bernoulli feedback on the output). We choose to frame our decision environment to increase subjects' understanding of our relatively complicated decision environment Alekseev, Charness and Gneezy (2017).

is then assigned an "IQ rank score" depending on their ranking within their assigned group of 19 other participants with random tie-breaking. Specifically, a subject that ranks within the i^{th} quintile is assigned an IQ rank score of $20i$. We then establish subjects' true ability parameters in the main part of the experiment as their IQ rank scores, i.e. $a_0 = 20i$. We incentivize subjects by paying them $25 if a randomly selected answer in the IQ test is correct.

3.2 Part 2: Establishing Mental Models

The main goal of this part of the experiment is to elicit prior beliefs on own ability that we use to establish mental models in the third part of the experiment. In order to achieve this goal, we ask subjects how they think they rank in their randomly constructed group of 20 people based on their IQ test scores. The reason we choose to measure relative overconfidence (or "overplacement") rather than absolute overconfidence (or "overestimation") is that previous experiments find a greater scope for overconfidence when it is measured in relative terms.[3] Another important design choice is that we ask our subjects to state their beliefs over quintiles rather than the more conventional way of measuring overconfidence using 2-quantiles. The reason we ask our subjects to state their full belief distribution over quintiles is that we want to provide scope for model misspecification while limiting the complexity of the belief elicitation procedure. The idea is that eliciting beliefs using smaller quantiles would create situations in which subjects predominantly assign positive probability to each quantile more frequently which would then render a majority of our subjects as correctly specified agents. On the other hand, if we elicit belief using larger quantiles, we complicate the belief elicitation procedure as subjects are required to state their full belief distribution over each quantile.

In order to simplify the belief elicitation procedure over quintiles, we use five sliders. Each quintile is associated with a slider. Subjects assign a total likelihood of 100% over five different quintiles through associated sliders at a precision of two decimal points. We use a standard incentive compatible mechanism to pay for the belief elicitation (Hossain and Okui (2013)). A critical design choice here is that we elicit beliefs over quintiles as full belief distribution. Eliciting full belief distribution with high precision allows us to sharply draw a line between subjects

[3] In particular, research in psychology documents that people "overplace" themselves in easy tasks (Larrick, Burson and Soll (2007), Moore and Small (2007)). We specifically choose the Raven matrices to benefit from this "easy" effect. Indeed, the average number of correct answers in our experiment is 6.78 out of 10.

with misspecified mental models and correctly specified mental models. In this regard, our belief elicitation procedure is a key element of the design as it allows us to strictly follow the theoretical conceptualization of overconfidence as a model misspecification.

3.3 Part 3: Learning Environment

After establishing an ability parameter for each subject and eliciting subjects' beliefs over their ability parameters, the only ingredient that is missing to construct a decision environment that is identical to our theoretical framework is the assignment of a fundamental to each subject. At the beginning of this part of the experiment, we randomly draw a fundamental for each subject using a discrete uniform distribution. The assignment of fundamental being independent from subjects' ability parameters is clearly communicated to subjects. Once subjects are assigned fundamentals, each subject faces an objective decision environment and they have mental models of their decision environments. As the researchers, we can observe both the objective and mental models of each subject.

3.3.1 Feedback Design and Minimizing Problems Related to Bayesianism

A crucial part of the experiment is the feedback that we provide to our subjects. Since our predictions are valid under Bayesian learning and the decision environment of our experiment is fairly complex, we help our subjects substantially make accurate inferences using the feedback regarding the fundamental. Consistent with the idea that people's learning about non-ego relevant variables is more in line with Bayesianism compared to learning about ego-relevant variables, we completely rule out that possibility that subjects' learning about the fundamental is inconsistent with Bayesian learning. We implement this critical feature of the design by providing subjects with a simple report that we frame as "the Statistician's Report" where we show subjects the Bayesian posterior mean of the fundamental conditional on each ability level. The report is updated in every period based on the feedback generated by the subject up until that period. Figure 1 presents an example of these reports.

Figure 1: The Statistician's Report

The Statistician's Report	
Your IQ Rank Score	**Project Quality**
20	75
40	71
60	69
80	66
100	65

Notes: Figure shows an example report. Each row shows how a particular IQ rank score corresponds to an expected project quality where the expectation is taken over the Bayesian posterior conditional on the IQ rank score.

3.3.2 Endogeneity of Feedback across Treatments

The only difference between our treatments **Exogenous** and **Endogenous** is if the feedback we provide to subjects is exogenous or endogenous to their actions. In treatment **Exogenous**, we provide subjects feedback based on the highest possible action,which is 100, not their actual actions. In contrast, we provide subjects feedback based on their actual actions in treatment **Endogenous**. We frame the lack of endogeneity of feedback to actions in treatment **Exogenous** as the company not being able to implement the subjects' recommended investment decisions immediately and instead investing an originally planned investment amount of 100 throughout the experiment. On the other hand, we frame the endogeneity of feedback in treatment **Endogenous** as the company implementing subjects' investment recommendations immediately instead of implementing their originally planned investment amount of 100. Note that we still mention the fact that there is an originally planned investment amount of 100 in treatment **Endogenous** to control for the potential anchoring effects.

The reason we choose the fixed action in treatment **Exogenous** as the highest possible action for each subject is two-fold. First, the predictions on self-defeating learning requires an exogenous action that is at or above the fundamental in our environment. Choosing the maximum possible action ensures that this requirement is satisfied irrespective of the realization of the fundamental. Second, the difference between predictions in **Exogenous** and **Endogenous** treatments in terms of expected fundamentals and chosen actions increases with the fixed action chosen in **Exogenous** treatment. Thus, choosing the highest possible action as the fixed action generates the largest possible treatment effect in theory.

3.3.3 Myopically Optimal Actions and the Calculator

A crucial element of the literature on learning with misspecified models is that agents take optimal actions (myopic or dynamic) in each period using their subjective expectations on the decision variables. In order to create an environment that allows our subjects to easily take myopically optimal actions, we choose a strictly concave output function that has a unique and simple optimal decision rule that only depends on the fundamental: "match your action to your expectation of the fundamental." We communicate this simple optimal action rule to our subjects as well as going through the details of how subjects can arrive at this conclusion on their own. We further test subjects' understanding of the optimal action rule through understanding quizzes.

We go a step further to make it even easier for our subjects to take myopically optimal actions by providing them with a calculator that takes their beliefs on their ability parameter as input and produces the myopically optimal action for that period as output. Hence for any subjective belief the subjects may have on their own ability parameter, they can accordingly calculate a myopically optimal action.

We attach the calculator to the Statistician's Report and ask subjects to enter their beliefs about their IQ rank score in corresponding rows. Once subjects enter their beliefs, the calculator produces the corresponding myopically optimal action using the Statistician's Report. More specifically, the calculator calculates a myopically optimal action by taking a weighted average of the expected project qualities with weights coming from the subjects' assigned likelihoods on each IQ rank score. Figure 2 provides an example.

Figure 2: Calculator

Calculator	The Statistician's Report	
Enter Likelihood (out of 100)	Your IQ Rank Score	Project Quality
	20	50
	40	50
	60	50
	80	50
	100	50
CALCULATE		

Notes: Figure shows an example report with the calculator attached to it. Subject are asked to enter their beliefs about their IQ rank score in corresponding rows as input. The calculator then calculates a myopically optimal action by taking a weighted average of the expected project qualities with weights coming from the subjects' assigned likelihoods on each IQ rank score.

We choose to help subjects calculate myopically optimal actions rather than dynamically optimal ones for several reasons. The first and main reason is that myopically optimal actions are significantly easier to explain to our subjects. Second, investigating how people learn under myopically optimal actions is of empirical relevance as previous research documents many instances where people narrowly bracket their decisions. Third, assuming that subjects' actions are consistent with myopic optimization allows us to interpret their choices reflecting their mean beliefs on the fundamental in our **Endogenous** treatment.[4]

Once a subject provides an input to the calculator and calculates an optimal action, a decision box appears where subjects are allowed to submit their actions. Figure 3 provides an example.

[4]Note that subjects' actions in treatment **Exogenous** directly correspond to their expectation of the fundamental as the difference between myopic and dynamically optimal actions vanishes due to the fact that there is no scope for experimentation in treatment **Exogenous**, that is, subjects sample from the same distribution throughout the experiment irrespective of their actions.

Figure 3: Decision Screen - Period 1

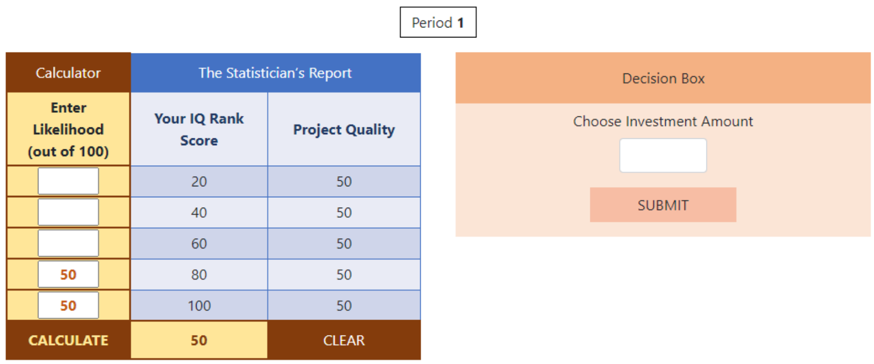

Notes: Figure shows an example decision screen. Once a subject provides an input to the calculator and calculates an optimal action, a decision box appears where subjects are allowed to submit their actions.

Although we require subjects to enter their beliefs about their IQ rank scores and calculate a myopically optimal action in each period, we still allow our subjects to submit actions that are distinct from what they obtain from the calculator as they might appreciate actions that are higher than the myopically optimal ones due to their informational value. The reason we require subjects to use the calculator in each period is to get a sense of the evolution of their beliefs on their own ability throughout the experiment.

3.3.4 The Amount of Opportunities to Learn

Since our goal is to investigate how subjects with misspecified models learn about an external fundamental, we design the experiment so that subjects have plenty of opportunities to take actions, generate feedback and learn from the feedback that they generate. We create three connected subparts for this part of the experiment. Although the number of actions subjects take are similar in each subpart, the amount of feedback that is generated through the implemented action gradually increases.

Subpart 1: Periods 1 to 10

Subjects start this part of the experiment by taking 10 actions. After each action, subjects get a binary feedback on the implemented action.

Subpart 2: Periods 11 to 100

Starting from the 11th period up to 100th period, subjects take actions every 10 periods, that is, in periods 11,21,31,...,91. We count the actions subjects take in each of these periods towards the following 9 periods and provide aggregate feedback for every 10 periods. For instance, when a subject takes an action in period 11, the same action also counts as the action the subject takes for periods 12 to 20. The subject is then provided aggregate feedback on implemented actions from periods 11 to 20. Subjects take a total of 9 actions in this subpart and get feedback from 90 periods.

Subpart 3: Periods 101 to 1000

When subjects reach the 101st period, they start taking actions every 100 periods until period 1000, that is, in periods 101,201,301,...,901. Similar to the previous subpart, we count the actions subjects take in each of these periods towards the following 99 periods and provide aggregate feedback for every 100 periods. For instance, when a subject takes an action in period 101, the same action also counts as the action the subject takes for periods 102 to 200. The subject is then provided aggregate feedback on implemented actions from periods 101 to 200. Subjects take a total of 9 actions in this subpart and get feedback from 900 periods.

Subjects continuously move from the first subpart to the third subpart and are informed of the beginning of a new subpart along the way. The subpart structure we implement follows from earlier designs carefully studying learning (Esponda, Vespa and Yuksel (2020)) and allows us to generate a significant amount of feedback without increasing the duration of the experiment.

3.3.5 Subjects' Payments

We incentivize our subjects by paying them a fixed reward of $25 if the feedback in a randomly chosen period is positive in treatment **Endogenous**. In treatment **Exogenous**, we re-draw a feedback for each period that is generated through subjects' actions in the experiment. This creates an incentive compatible mechanism for subjects to take optimal actions in a manner that is equivalent to binarized scoring rule Hossain and Okui (2013).

3.4 Part 4: Re-Examining Mental Models

In the fourth part of the experiment, we elicit subjects' beliefs about their ability parameter for a second time. Eliciting the full posterior belief distribution after subjects receive 1000 periods worth of feedback on their own ability allows us to answer if subjects retain their initial mental models or switch to alternative models.

The belief elicitation procedure is identical to the second part of the experiment. Subjects use five sliders to indicate their beliefs about which quintile their rank in their randomly constructed group of 20 people based on their IQ test scores. We use binarized scoring rule to incentivize subjects to truthfully report their beliefs about their IQ rank score.

3.5 Part 5: Exit Survey

We finalize the experiment by asking subjects control questions about their gender, their year of study, if they are enrolled in a STEM major and if they have taken a college-level statistics class.

3.6 Procedural Details

We conducted our experiment online using the subject pool of UCSB Experimental and Behavioral Economics Laboratory. The experiment was coded using o-Tree software (Chen, Schonger and Wickens (2016)). A total of 128 subjects, recruited through ORSEE (Online Recruitment System For Economic Experiments) ?. The average payment per subject was $27.6 including a $10 show-up fee. Each session lasted for 105 minutes.

4 Results

4.1 Identifying Misspecified Mental Models

We define an overconfident agent as one whose prior belief assigns zero mass on their true and all lower level IQ rank scores. Our clear identification strategy directly follows from the theoretical conceptualization of overconfidence as a misspecified mental model. Note that our conceptualization of overconfidence is more stringent than the typical conceptualization of overconfidence as having mean or median beliefs laying above the actual "ability" parameter. Figures 4 and 5 respectively display examples of overconfident subjects and subjects with correctly specified mental models.

Figure 4: Overconfidence as a Misspecified Mental Model

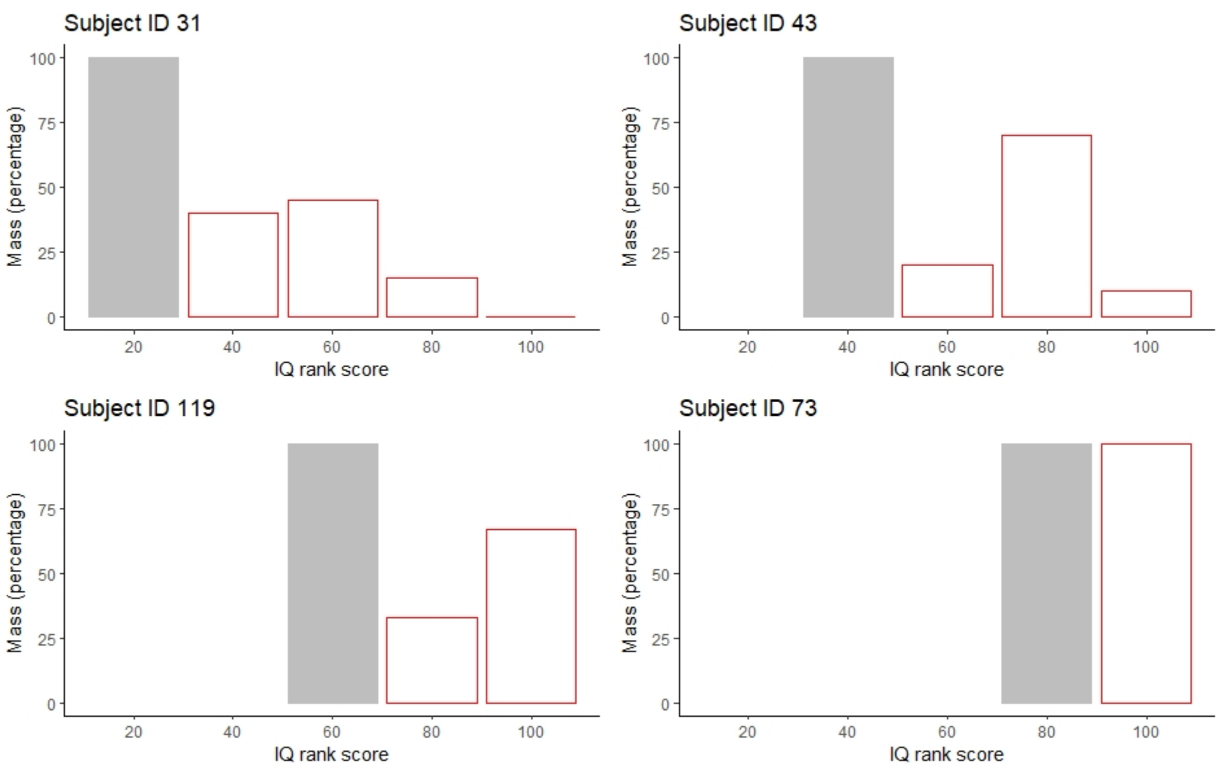

Notes: Figure shows the prior beliefs of selected subjects on their IQ rank scores. The grey bars display subjects' true IQ rank scores. The red bars display subjects' priors as probability mass functions.

We identify a total of 42 overconfident subjects (out of 124) distributed almost evenly across our two treatments **Exogenous** and **Endogenous**. The share of overconfident subjects in treatment **Exogenous** is 31.25% whereas the share of overconfident subjects in treatment **Endogenous** is 34.4%. The fact that we generate a significant amount of model misspecification through a simple task lends support to the main premise of the theoretical literature that people might form priors that sharply exclude the possibility of truth. We do not identify underconfidence as a model misspecification in our data, all remaining subjects in our experiment assigns some positive mass on their true IQ rank score.

Figure 5: Correctly Specified Mental Models

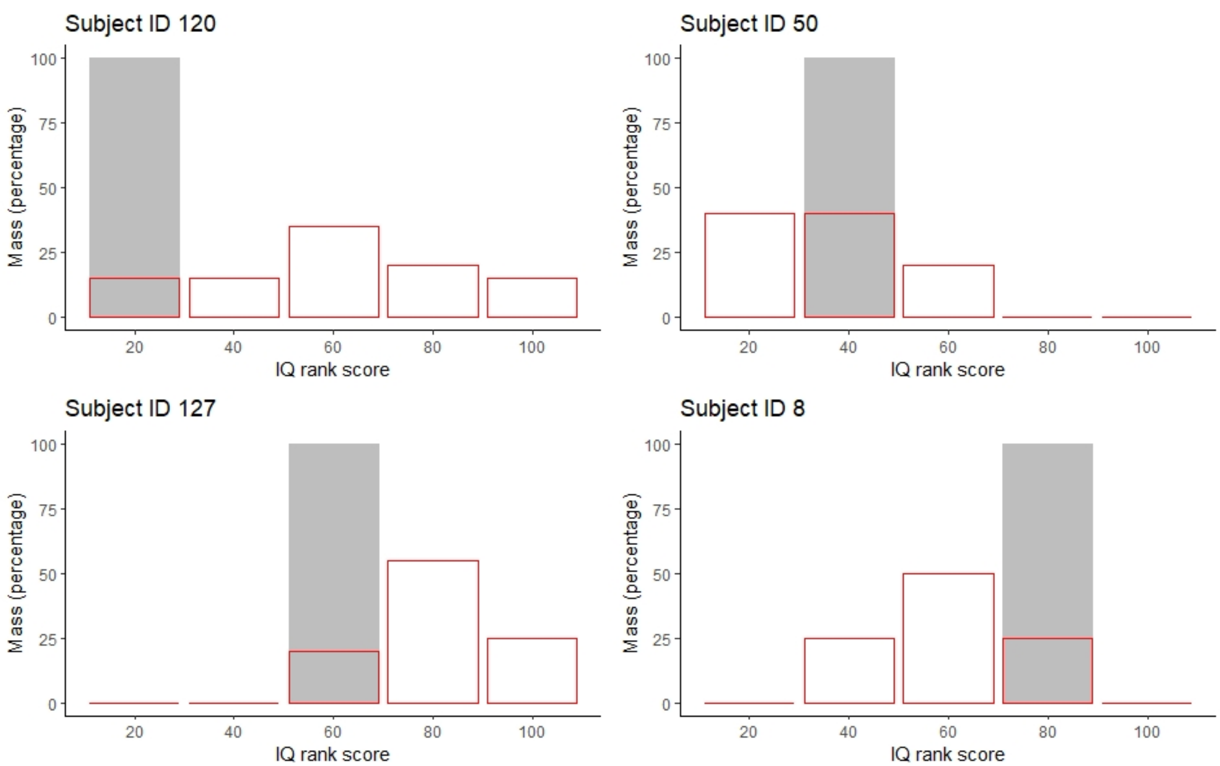

Notes: Figure shows the prior beliefs of selected subjects on their IQ rank scores. The grey bars display subjects' true IQ rank scores. The red bars display subjects' priors as probability mass functions.

4.2 Do Misspecified Mental Models Generate Suboptimal Behavior?

In this subsection, we aim to answer two questions. First, we ask if overconfident subjects' learning processes lead them to take suboptimal actions. Second, we ask how overconfident and correctly specified subjects' actions compare against the myopically optimizing Bayesian benchmark.

A First Look at How Overconfidence Generates Growingly Suboptimal Behavior

Figure 6 shows the evolution of actions for overconfident and correctly specified subjects when we aggregate the data from treatments **Exogenous** and **Endogenous**. Since subjects' fundamentals are drawn uniform randomly over integers from 0 to 100, we expect the fundamentals and hence the first-best optimal actions to average around 50. Indeed, Panel A of Figure 6 shows that average first-best optimal action, that is the action a subject would take if they were to know the true fundamental, for overconfident subjects is 53 and the average first-best optimal action for correctly specified subjects is 50.73. We find that correctly specified subjects' learning process do not lead them away from the first-best optimal action, the average action for correctly specified agents remain around the first-best optimal action throughout the experiment and moves closer to the first-best optimal action starting from period 11.

The behavior of overconfident subjects is dramatically different from correctly specified subjects. We find that overconfident subjects start out by taking higher actions than the first-best optimal actions reflecting their optimistic beliefs about the fundamental. After 10 periods, we find overconfident subjects start taking actions that are significantly lower than the first-best optimal action, reflecting their pessimistic beliefs about the fundamental and persistently keep doing so for the remainder of the experiment. By Period 901, we find a clear difference in behavior among overconfident and correctly specified subjects: overconfident subjects are on average taking actions that are significantly lower than optimal whereas correctly specified subjects on average are taking optimal actions.

Panel B of Figure 6 shows that the stark difference in behavior we observe is consistent with the theoretical predictions. When we simulate Bayesian learning with myopically optimal actions for each subject taking their mental models as given, we find that correctly specified subjects should take an average action of 49.05 in Period 901, whereas overconfident subjects should take an average action of 36.37. The difference in simulated average actions between correctly specified and overconfident subjects is significant ($p < 0.01$). We find that correctly specified agents take actions that are consistent with Bayesian learning. On the other hand, although overconfident subjects' actions are moving towards the simulated Bayesian action, the average action is still far away from the simulated Bayesian action.

Figure 6: Evolution of Actions

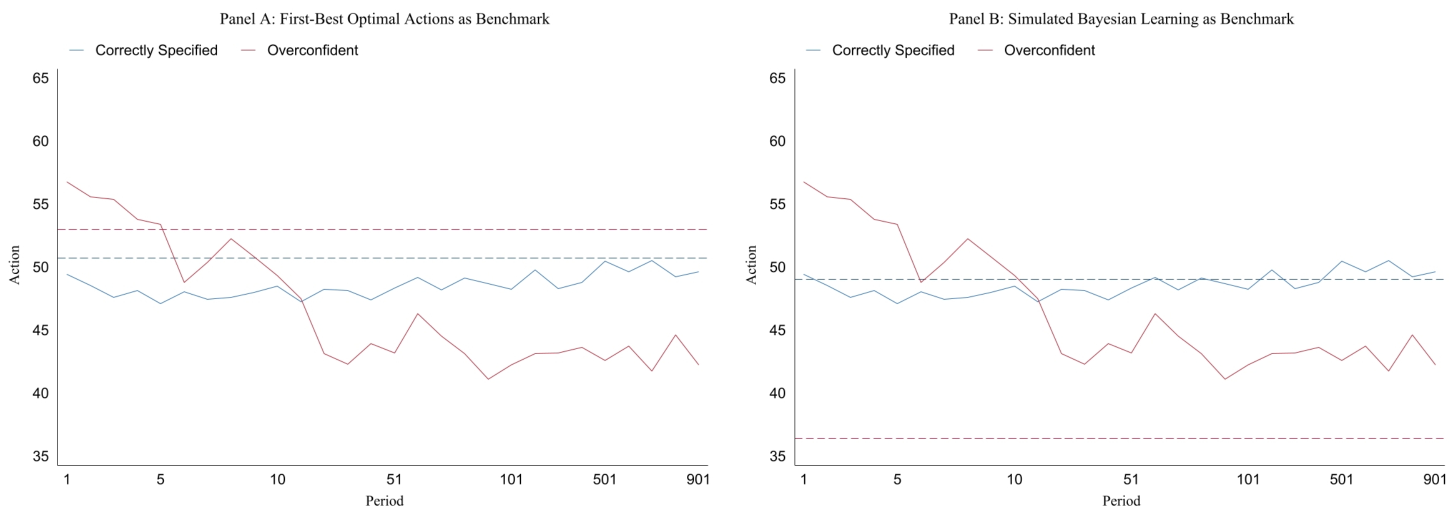

Notes: Both Panel A and Panel B show the average action for correctly specified and overconfident subjects across periods. Panel A presents the first-best optimal action as a benchmark. The blue dashed line in Panel A represents the average first-best optimal action for correctly specified subjects. The red dashed line in Panel A represents the average first-best optimal action for overconfident subjects. Panel B presents the average action a myopically optimizing Bayesian agent would take in the last period of the experiment as a benchmark. The simulations are conducted using each subject's prior beliefs about their abilities. The blue dashed line in Panel B represents the average simulated Bayesian action for correctly specified subjects whereas the red dashed line in Panel B represents the average simulated Bayesian action for overconfident subjects.

Statistical Differences in Learning with and without a Misspecified Model

We use a displacement measure $\Delta_{OPT} = e - e^*(\phi_0)$ to capture how far each action e is relative to the first-best optimal action. Note that this measure is positive for actions that are greater than the first-best optimal action and negative for actions that are smaller than the first-best optimal action. We then estimate the following displacement-from-benchmark regression in different periods: $\Delta_{OPT} = \alpha + \beta M + \varepsilon$ where α captures displacement from first-best optimal action for correctly specified subjects, M is a dummy variables that takes the value 1 for overconfident subjects, β captures the difference in displacement-from-benchmark for overconfident subjects and ε is an error term.

Table 1 presents the estimation results. First, note that correctly specified subjects do not significantly displace themselves from the first-best optimal action neither at the beginning nor towards the end of the experiment. On the other hand, we find that overconfident subjects start out positively yet not significantly displacing themselves from the first-best optimal action ($p = 0.48$). However, we find that initial positive displacement of overconfident subjects turn significantly and persistently negative in the later periods of the experiment. In particular, throughout the last 400 periods of the experiment, overconfident subjects displace themselves around 10 points away from the first-best optimal action. Figure A1 in Appendix A presents these patterns in detail.

Table 1: Estimation of the Effect of Overconfidence on Displacement Relative to the First-Best Optimal Action

	Dependent Variable: Δ_{OPT}			
	(1)	(2)	(3)	(4)
β	5.087	-10.13**	-11.03***	-9.662**
	(6.348)	(3.231)	(3.032)	(2.887)
α	-1.326	-0.293	-0.217	-1.113
	(3.556)	(2.184)	(2.195)	(1.746)
Observations	128	128	128	128
Period	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for correctly specified and overconfident agents. Each column conducts the estimation $\Delta_{OPT} = \alpha + \beta M + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action. The observations within periods aggregated across treatments **Exogenous** and **Endogenous**.

We use a second displacement measure $\Delta_{BAYES} = e - e^*(E_{\Pi_{sim}}[\phi])$, where Π_{sim} is the simulated posterior distribution on ϕ in the last period of the experiment, to capture how far each action e

is relative to the simulated Bayesian action for the last period in the experiment.[5][6] We separately estimate a displacement-from-benchmark regression for correctly specified and overconfident subjects in different periods: $\Delta_{BAYES} = \alpha + \varepsilon$ where α captures displacement from the simulated Bayesian action for either correctly specified or overconfident subjects, and ε is an error term.

Table 2: Estimation of Displacement Relative to the Simulated Bayesian Action

	Panel A: Correctly Specified				Panel B: Overconfident			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
α	0.357	1.389	1.465	0.570	20.40***	6.208***	5.389**	5.861*
	(3.427)	(1.955)	(1.965)	(1.475)	(4.296)	(1.687)	(1.567)	(2.214)
Observations	86	86	86	86	42	42	42	42
Period	1	501	701	901	1	501	701	901

Notes: The table presents the average displacement from the simulated Bayesian action relative to the last period for correctly specified and overconfident agents. Each column conducts the estimation $\Delta_{BAYES} = \alpha + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action. The observations within periods are aggregated across treatments **Exogenous** and **Endogenous**.

Table 2 presents the estimation results. Note that correctly specified subjects on average do not displace themselves from the simulated Bayesian action. On the other hand, overconfident subjects are systematically over the simulated Bayesian action. However, we see that the extent of displacement for the overconfident subjects is getting smaller as subjects move along the experiment's time horizon. Figure A2 in Appendix A presents these patterns in detail.

We summarize our findings in this subsection in the following result:

Result 1. *Overconfident subjects' learning processes lead them to take growingly suboptimal actions throughout the experiment's time horizon. On the other hand, correctly specified subjects' learning processes do not generate a systematic deviation from the first-best optimal action throughout the experiment. Moreover, overconfident subjects' learning processes yield outcomes that are* ***less***

[5] To be more precise, $e^*(E_{\Pi_{sim}}[\phi])$ is the action that a myopically optimizing Bayesian agent would take in the last period of the experiment.

[6] Similarly, this measure is positive for actions that are greater than the simulated Bayesian action and negative for actions that are smaller than the simulated Bayesian action.

suboptimal than the Bayesian prediction while correctly specified subjects' learning process is fully consistent with the Bayesian prediction.

4.3 Does Endogenous Learning Exacerbate Suboptimal Behavior?

In this subsection, we aim to answer if endogenous learning exacerbate suboptimal behavior when the feedback subjects receive is endogenous to their actions. Theoretical predictions are such that endogenous learning should exacerbate overconfident agents' suboptimal behavior. For correctly specified agents, the theory's predictions are ambiguous. However, using simulations, we find that endogenous learning should not lead to a change in behavior for correctly specified subjects by the end of the experiment. We start this subsection by comparing the behavior of overconfident subjects in **Exogenous** and **Endogenous**. We then turn to correctly specified subjects and compare their behavior across treatments.

Behavior of Overconfident Subjects

Figure 7 shows the evolution of displacement relative to the first-best optimal action for overconfident subjects in **Exogenous** and **Endogenous**. We find that overconfident subjects in both treatments start out with actions that are close to the first-best optimal action. In both treatments, subjects exhibit negative displacement over time and we find that overconfident subjects in **Endogenous** exhibit greater negative displacement starting with Period 10. However, the difference in negative displacement vanishes towards the end of the experiment. Panel B provides further insights as to why we see the difference between the treatments vanish towards the end of the experiment. Overconfident subjects in **Exogenous** persistently move closer to the Bayesian prediction and meets the Bayesian prediction in the final period of the experiment whereas overconfident subjects in **Endogenous** decelerate their move towards the Bayesian prediction towards the end of the experiment.

Figure 7: Evolution of Displacement for Overconfident Subjects

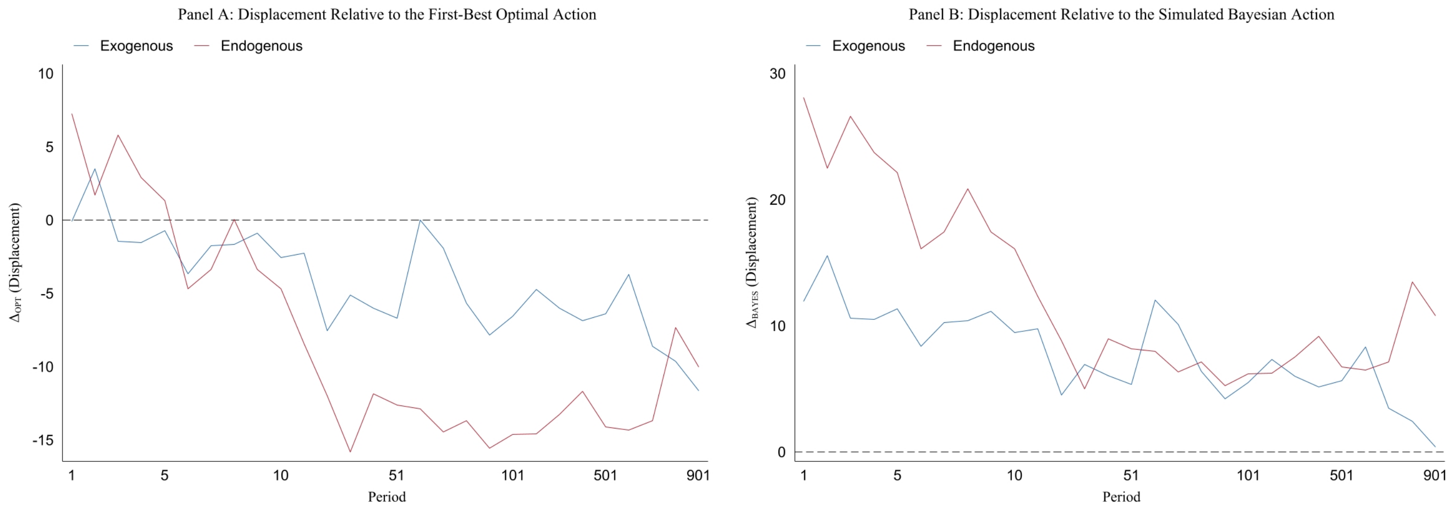

Notes: Panel A shows the average displacement relative to the first-best optimal action for overconfident subjects in treatment **Exogenous** and **Endogenous** across periods. Panel B shows the average displacement relative to the simulated Bayesian action for overconfident subjects in treatment **Exogenous** and **Endogenous** across periods. Each observation in a period corresponds to an individual action.

Table 3 provides the estimates of the treatment effect using displacement relative to the first-best optimal action as a benchmark. We estimate the regression $\Delta_{OPT} = \alpha + \theta T + \varepsilon$ where α captures the average displacement-from-benchmark for subjects in **Exogenous**, T is a dummy variable that takes the value 1 for subjects in **Endogenous**, θ captures the treatment effect and ε is an error term. The estimates for α across periods clearly show that overconfident subjects' actions in **Exogenous** growingly and significantly moves away from to the first-best optimal action exhibiting negative displacement from Period 1 to 901. On the other hand, the estimates for θ across periods show that there is no significant exacerbation of displacement for overconfident subjects in **Endogenous**. Table A1 provides statistical evidence that overconfident subjects in Exogenous act consistent with the Bayesian prediction by Period 901, while the behavior of overconfident subjects remain markedly different from the Bayesian prediction by Period 901.

Table 3: Estimation of the Treatment Effect for Overconfident Subjects

	Dependent Variable: Δ_{OPT}			
	(1)	(2)	(3)	(4)
θ	7.373	-7.683	-5.076	1.627
	(10.51)	(4.633)	(4.127)	(4.677)
α	-0.100	-6.404*	-8.588**	-11.63**
	(6.604)	(2.889)	(2.564)	(3.410)
Observations	42	42	42	42
Period	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for overconfident agents. Each column conducts the estimation $\Delta_{OPT} = \alpha + \beta T + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action.

Behavior of Correctly Specified Subjects

Panel A of Figure 8 displays the evolution of displacement relative to the first-best optimal action for overconfident subjects in **Exogenous** and **Endogenous**. There is no discernible difference between the actions of correctly specified subjects in our treatments. We see that in both treatments behavior remains close to the first-best optimal benchmark throughout the experiment. Panel B of Figure 8 displays the evolution of displacement relative to the simulated Bayesian actions. Again, we do not see any systematic deviation from the Bayesian benchmark for the duration of the experiment.

Figure 8: Evolution of Displacement for Correctly Specified Subjects

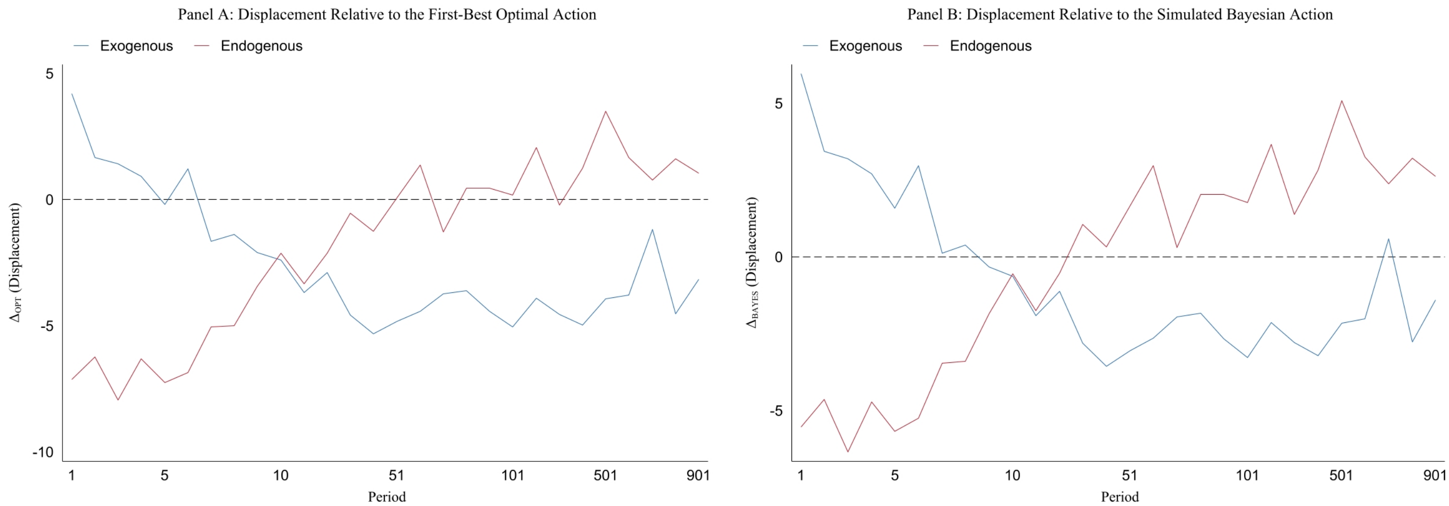

Notes: Panel A shows the average displacement relative to the first-best optimal action for correctly specified subjects in treatment **Exogenous** and **Endogenous** across periods. Panel B shows the average displacement relative to the simulated Bayesian action for overconfident subjects in treatment **Exogenous** and **Endogenous** across periods. Each observation in a period corresponds to an individual action.

Table 4 provides the estimates of the treatment effect using displacement relative to the first-best optimal action as a benchmark. We estimate the identical displacement-from-benchmark regression $\Delta_{OPT} = \alpha + \theta T + \varepsilon$ for correctly specified subjects. The estimates for α show that subjects in Exogenous exhibit negative displacement towards the end of the experiment although the magnitude of this move in each period is insignificant. The estimates for θ in each period indicate that subjects in **Endogenous** do not take significantly different actions compared to the subjects in **Exogenous**. Table A2 further documents that the subject behavior in both **Exogenous** and **Endogenous** are consistent with the simulated Bayesian action by the end of the experiment.

Table 4: Estimation of the Treatment Effect for Correctly Specified Subjects

	Dependent Variable: Δ_{OPT}			
	(1)	(2)	(3)	(4)
θ	-11.32	7.426	1.965	4.209
	(7.025)	(4.298)	(4.370)	(3.462)
α	4.205	-3.920	-1.177	-3.168
	(5.038)	(3.212)	(3.535)	(2.626)
Observations	86	86	86	86
Period	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for correctly specified subjects. Each column conducts the estimation $\Delta_{OPT} = \alpha + \theta T + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action.

We summarize our findings from this subsection in the following result:

Result 2. *Contrary to the theoretical prediction, endogenous learning does not exacerbate the extent of suboptimal behavior for overconfident subjects. Similarly, although now consistent with the theory, we do not detect a change in behavior for correctly specified subjects when feedback is endogenous to their actions. Moreover, we find that overconfident subjects' behavior deviates from the Bayesian benchmark when feedback is endogenous but not when feedback is exogenous. On the other hand, correctly specified subjects behavior do not deviate from the Bayesian benchmark irrespective of the endogeneity of feedback.*

4.4 Learning About One's Self

We have so far investigated how subjects learn about the external decision variable in their environment. In this subsection, we turn to how subjects learn about their "ability" parameters. A Berk-Nash equilibrium of the single agent problem we investigate is one in which beliefs about the ability parameter concentrate on the prior where there is no self-learning as we have discussed earlier. Simulating a Bayesian learning model for the finite duration of our experiment yields posteriors consistent with this equilibrium by the end of the experiment, that is Period 1000. Simulations also confirm that there should be no self-learning for the duration of the experiment in subjects in **Exogenous**. We start this subsection by first looking at how overconfident subjects learn about their ability compared to correctly specified subjects at the aggregate. We then look at how endogenous learning affects learning about self for overconfident and correctly specified subjects.

Throughout this subsection, we use a displacement measure $\Delta_a = E_p[a] - a_0$ to capture how far each expected ability level $E_p[a]$ is relative to the true ability a_0 where the expectation is taken using the probability mass function p on a. Note that this measure is positive for beliefs that generate an expected ability level that is greater than the true ability level.

Differences in Self-Learning between Overconfident and Correctly Specified Subjects

Figure 9 presents a comparison of the displacement of expected abilities relative to the true ability using three different probability mass functions: subject's elicited *prior* on their ability, subject's elicited *posterior* on their ability and the simulated *Bayesian posterior* for the subject. Panel A shows that overconfident agents' expectations of their abilities move towards their true abilities after receiving 1000 periods worth of feedback. This move is significant at conventional levels (p = 0.02). The significant reduction in displacement relative to the true ability is also inconsistent with Bayesian learning. We find that subjects posterior means average 5.96 lower than the simulated Bayesian posteriors (p = 0.04). Panel B documents that correctly specified agents' posterior expectations about their abilities do not significantly differ from their prior expectations (p = 0.16) or the Bayesian benchmark (p = 0.21).

Figure 9: Displacement of Expected Ability Relative to the True Ability

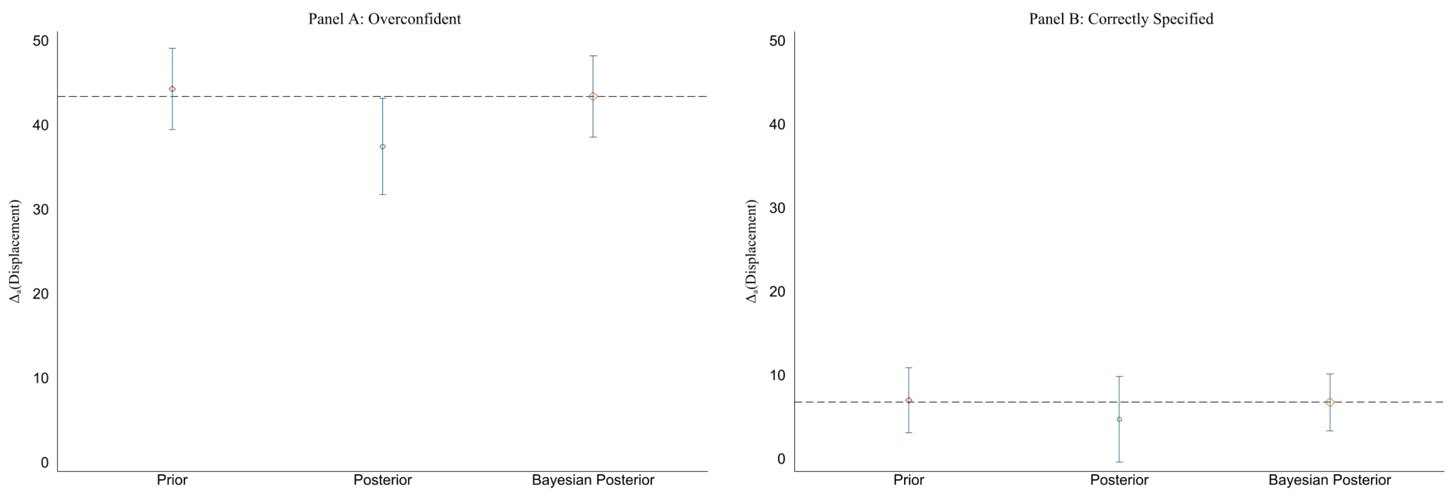

Notes: Figure shows the average displacement of expected ability levels relative to the true abilities using subjects' priors, posteriors, and simulated Bayesian posteriors. Panel A focuses on overconfident subjects where as Panel B focuses on correctly specified subjects. The black dashed line indicates the Bayesian benchmark. Whiskers indicate 95% confidence intervals.

One might then be curious if overconfident subjects move their expectations on their abilities towards their true abilities within the confines of their initial mental models. Table 5 provides evidence on how subjects' learning processes may lead them to completely switch their mental models. Although the majority of overconfident subjects stick with their initial mental models, we find that 22% of overconfident subjects end up assigning some probability to their true ability level after receiving feedback for 1000 periods. On the other hand, 13% of correctly specified subjects end up assigning no probability their true ability level at the end of their learning process.

Table 5: Switching Mental Models

		Posterior Models		
		Overconfident	Correctly Specified	Underconfident
Prior Models	Overconfident	78%	22%	0%
	Correctly Specified	13%	85%	2%

How Does Endogeneity of Feedback Affect Self-Learning for Overconfident Subjects?

According to the Bayesian benchmark, there is no difference in self-learning depending on the endogeneity of feedback. We expect overconfident subjects to exhibit virtually no self-learning in both **Exogenous** and **Endogenous** after 1000 periods. Panel A of Figure 10 documents that subjects in **Exogenous** somewhat learn their true ability after 1000 periods as the expected posterior beliefs show a smaller displacement from the true ability. The difference in prior and posterior means is insignificant ($p = 0.10$) and subjects' posterior mean is consistent with Bayesian posterior mean ($p = 0.17$). On the other hand, Panel B documents that posterior means of subjects in **Endogenous** considerably move towards their true abilities, yet the move is not significant at conventional levels ($p = 0.07$).

Figure 10: Displacement of Expected Ability Relative to the True Ability - Overconfident Subjects

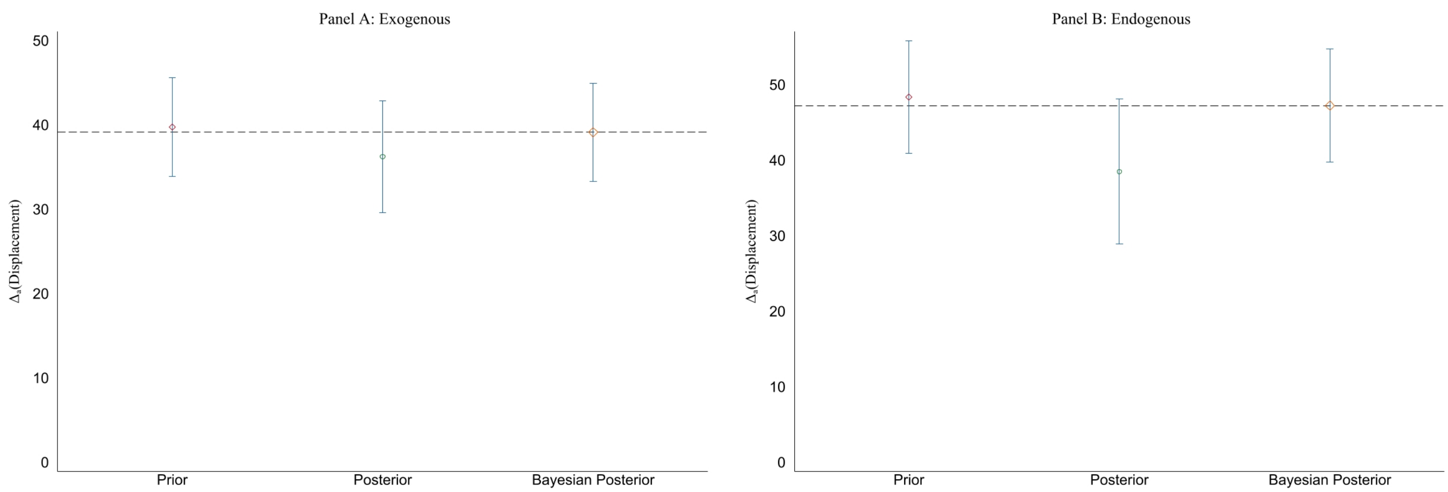

Notes: Figure shows the average displacement of expected ability levels relative to the true abilities using subjects' priors, posteriors, and simulated Bayesian posteriors. Panel A focuses on overconfident subjects in **Exogenous** where as Panel B focuses on overconfident subjects in **Endogenous**. The black dashed line indicates the Bayesian benchmark. Whiskers indicate 95% confidence intervals.

An important point that is worth emphasizing here is the increased self-learning with endogenous feedback is consistent with our earlier finding that endogenous learning does not exacerbate suboptimal behavior. If endogeneity of feedback leads subjects to better learn their own abilities, then subjects should take actions that are closer to the first-best optimal action in the main part of the experiment.

How Does Endogeneity of Feedback Affect Self-Learning for Correctly Specified Subjects?

As in the case of overconfidence, Bayesian learning does not predict subjects' self-learning to depend on the endogeneity of feedback. We expect no difference in mean prior and posterior beliefs for subjects in **Exogenous** and **Endogenous**. Panel A of Figure 11 confirms the Bayesian prediction for subjects in **Exogenous**: there is virtually no difference in prior and posterior means ($p = 0.92$). On the other hand, we find posterior means to move significantly closer to the subjects' true abilities in **Endogenous** ($p = 0.04$). The difference in posterior and the Bayesian posterior means in **Endogenous** is also significant ($p = 0.03$).

Figure 11: Displacement of Mean Beliefs Relative to the True Ability - Correctly Specified Subjects

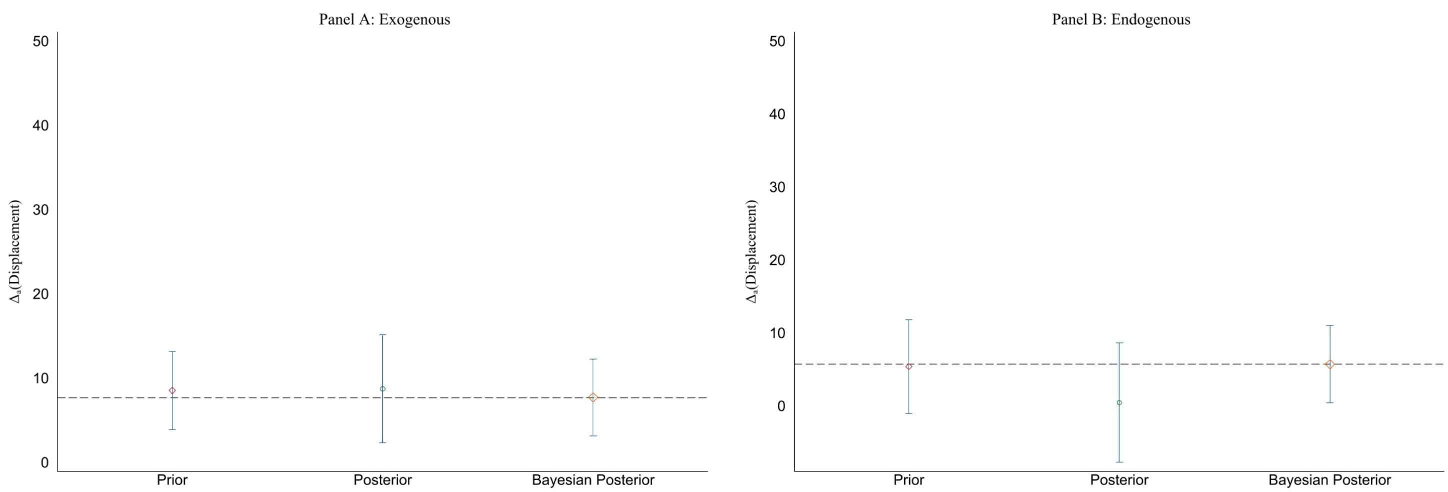

Notes: Figure shows the average displacement of expected ability levels relative to the true abilities using subjects' priors, posteriors, and simulated Bayesian posteriors. Panel A focuses on correctly specified subjects in **Exogenous** where as Panel B focuses on correctly specified subjects in **Endogenous**. The black dashed line indicates the Bayesian benchmark. Whiskers indicate 95% confidence intervals.

We summarize our findings from this subsection in the following result:

Result 3. *Inconsistent with Bayesian learning, overconfident subjects' posterior expectations about their ability significantly move closer to their true abilities. On the other hand, consistent with Bayesian learning, correctly specified subjects' prior and posterior expectations about their own ability do not differ. Moreover, both overconfident and correctly specified subjects exhibit greater self-learning when feedback is endogenous to their actions.*

5 Conclusion

In this paper we use people's tendency to hold optimistic beliefs about their abilities to generate model misspecification and investigate the implications of overconfidence as a misspecified mental model on learning about own ability and a fundamental. We find that overconfident subjects develop pessimistic beliefs about the fundamental and take growingly suboptimal actions. On the other hand, we find that endogenous feedback does not exacerbate the extent of suboptimal behavior: a result that is inconsistent with the theoretical prediction. When we look at how subjects learn about their own ability, we find that 1000 periods' worth of objective feedback lead some overconfident subjects to open their models to the possibility of truth. The "weakening" of mental models we observe is consistent with previous evidence. Complementing the nascent experimental literature on learning with misspecified mental models, we find that the "weakening" of mental models is more pronounced with endogenous feedback, explaining why endogenous feedback may not exacerbate the extent of suboptimal behavior.

References

Alekseev, Aleksandr, Gary Charness, and Uri Gneezy. 2017. "Experimental methods: When and why contextual instructions are important." *Journal of Economic Behavior & Organization*, 134: 48–59.

Barber, Brad M, and Terrance Odean. 2001. "Boys will be boys: Gender, overconfidence, and common stock investment." *The quarterly journal of economics*, 116(1): 261–292.

Bénabou, Roland, and Jean Tirole. 2002. "Self-confidence and personal motivation." *The quarterly journal of economics*, 117(3): 871–915.

Berk, Robert H. 1966. "Limiting behavior of posterior distributions when the model is incorrect." *The Annals of Mathematical Statistics*, 51–58.

Camerer, Colin, and Dan Lovallo. 1999. "Overconfidence and excess entry: An experimental approach." *American economic review*, 89(1): 306–318.

Chen, Daniel L, Martin Schonger, and Chris Wickens. 2016. "oTree—An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance*, 9: 88–97.

Chen, Si, and Hannah Schildberg-Hörisch. 2019. "Looking at the bright side: The motivational value of confidence." *European Economic Review*, 120: 103302.

Coutts, Alexander. 2019. "Good news and bad news are still news: Experimental evidence on belief updating." *Experimental Economics*, 22(2): 369–395.

Eil, David, and Justin M Rao. 2011. "The good news-bad news effect: asymmetric processing of objective information about yourself." *American Economic Journal: Microeconomics*, 3(2): 114–38.

Ertac, Seda. 2011. "Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback." *Journal of Economic Behavior & Organization*, 80(3): 532–545.

Esponda, Ignacio, and Demian Pouzo. 2016. "Berk–Nash equilibrium: A framework for modeling agents with misspecified models." *Econometrica*, 84(3): 1093–1130.

Esponda, Ignacio, Emanuel Vespa, and Sevgi Yuksel. 2020. "Mental models and learning: The case of base-rate neglect."

Handel, Benjamin, and Joshua Schwartzstein. 2018. "Frictions or Mental Gaps: What's Behind the Information We (Don't) Use and When Do We Care?" *Journal of Economic Perspectives*, 32(1): 155–78.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein. 2014. " Learning Through Noticing: Theory and Evidence from a Field Experiment *." *The Quarterly Journal of Economics*, 129(3): 1311–1353.

Heidhues, Paul, Botond Kőszegi, and Philipp Strack. 2018. "Unrealistic expectations and misguided learning." *Econometrica*, 86(4): 1159–1214.

Hossain, Tanjim, and Ryo Okui. 2013. "The binarized scoring rule." *Review of Economic Studies*, 80(3): 984–1001.

Köszegi, Botond. 2006. "Ego utility, overconfidence, and task choice." *Journal of the European Economic Association*, 4(4): 673–707.

Larrick, Richard P, Katherine A Burson, and Jack B Soll. 2007. "Social comparison and confidence: When thinking you're better than average predicts overconfidence (and when it does not)." *Organizational Behavior and Human Decision Processes*, 102(1): 76–94.

Malmendier, Ulrike, and Geoffrey Tate. 2005. "CEO overconfidence and corporate investment." *The journal of finance*, 60(6): 2661–2700.

Malmendier, Ulrike, and Geoffrey Tate. 2008. "Who makes acquisitions? CEO overconfidence and the market's reaction." *Journal of financial Economics*, 89(1): 20–43.

Möbius, Markus M, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat. 2014. "Managing self-confidence." *NBER Working paper*, 17014.

Moore, Don A, and Deborah A Small. 2007. "Error and bias in comparative judgment: on being both better and worse than we think we are." *Journal of personality and social psychology*, 92(6): 972.

Moore, Don A, and Paul J Healy. 2008. "The trouble with overconfidence." *Psychological review*, 115(2): 502.

Svenson, Ola. 1981. "Are we all less risky and more skillful than our fellow drivers?" *Acta psychologica*, 47(2): 143–148.

Weinstein, Neil D. 1980. "Unrealistic optimism about future life events." *Journal of personality and social psychology*, 39(5): 806.

Zimmermann, Florian. 2020. "The dynamics of motivated beliefs." *American Economic Review*, 110(2): 337–61.

Appendices

Appendix A Additional Results

Figure A1: Evolution of Displacement Relative to the First-Best Optimal Action

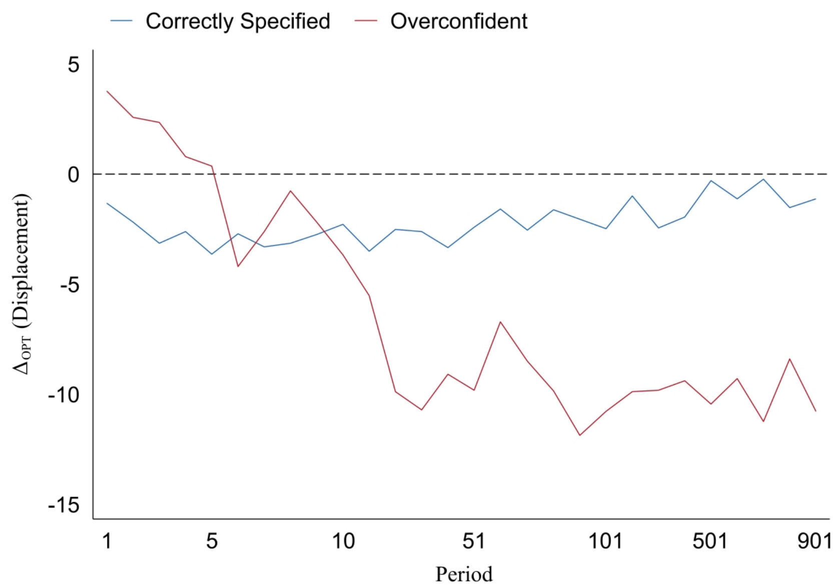

Notes: Figure shows the average displacement relative to the first-best optimal action for correctly specified and overconfident agents across periods. Each observation in a period corresponds to an individual action. The observations within a period are aggregated across treatments **Exogenous** and **Endogenous.**

Figure A2: Evolution of Displacement Relative to the Simulated Bayesian Action

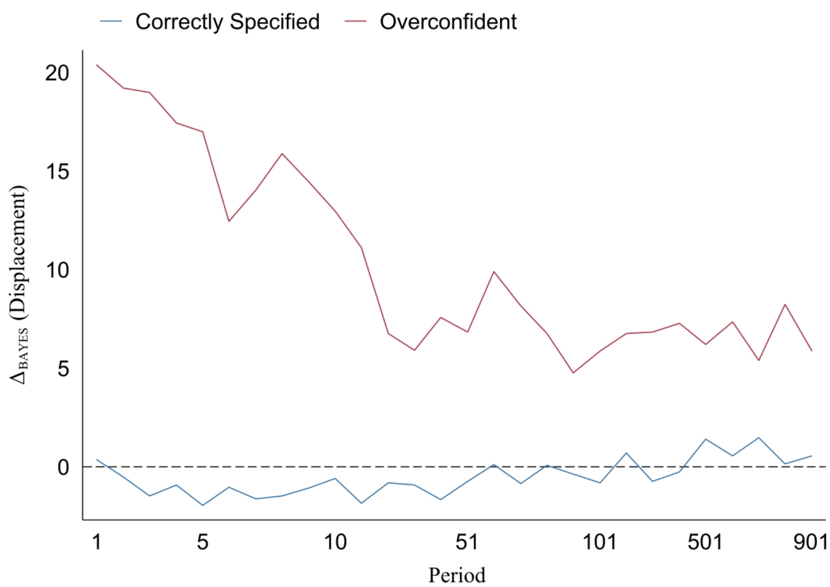

Notes: Figure shows the average displacement relative to the simulated Bayesian action for correctly specified and overconfident agents across periods. Each observation in a period corresponds to an individual action. The observations within a period are aggregated across treatments **Exogenous** and **Endogenous.**

Figure A3: Evolution of Actions for Overconfident Subjects

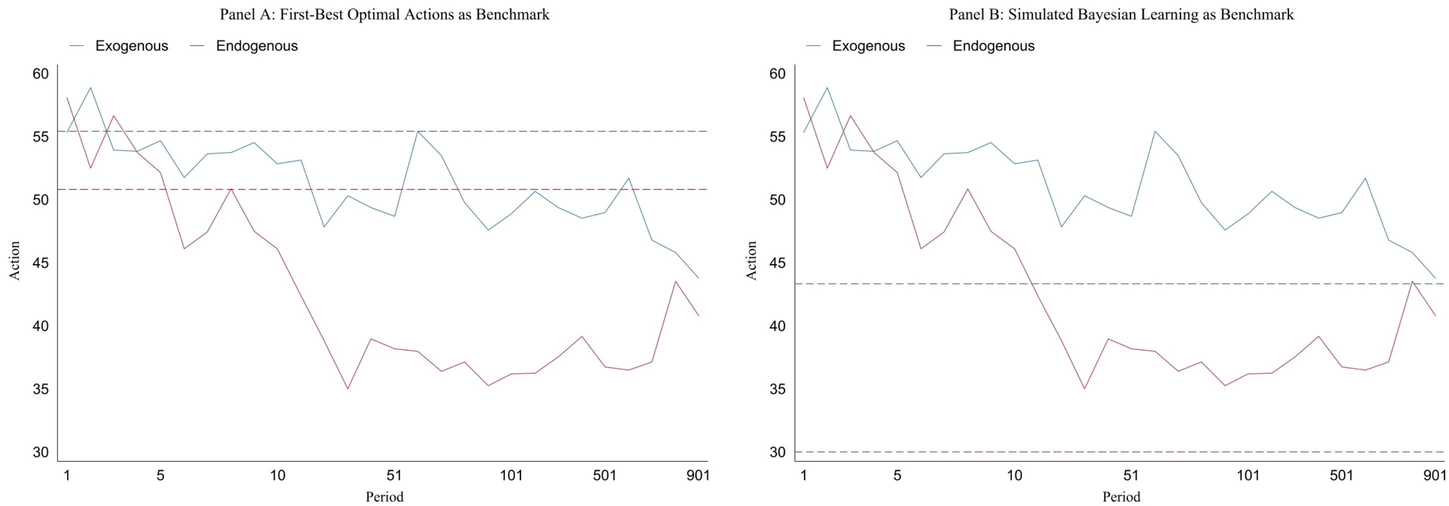

Notes: Both Panel A and Panel B show the average action separately for overconfident subjects in **Exogenous** and **Endogenous** across periods. Panel A presents the first-best optimal action as a benchmark. The blue dashed line in Panel A represents the average first-best optimal action for overconfident subjects in **Exogenous**. The red dashed line in Panel A represents the average first-best optimal action for overconfident subjects in **Endogenous**. Panel B presents the average action a myopically optimizing Bayesian agent would take in the last period of the experiment as a benchmark. The simulations are conducted using each subject's prior beliefs about their abilities. The blue dashed line in Panel B represents the average simulated Bayesian action for overconfident subjects in **Exogenous** whereas the red dashed line in Panel B represents the average simulated Bayesian action for overconfident subjects in **Endogenous**.

Figure A4: Evolution of Actions for Correctly Specified Subjects

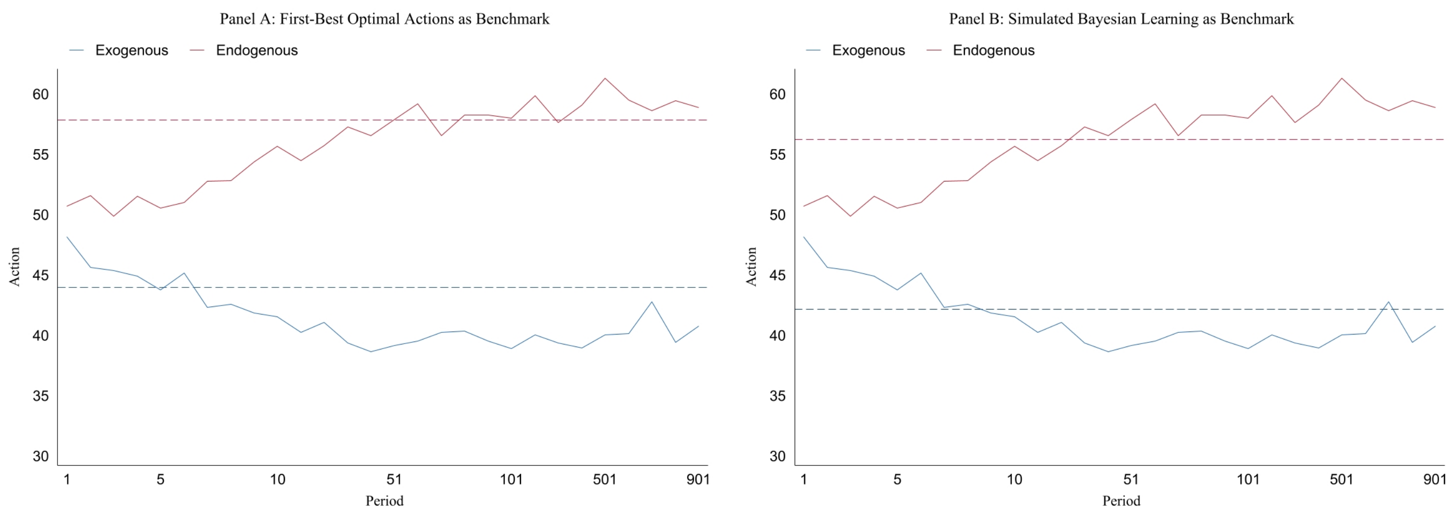

Notes: Both Panel A and Panel B show the average action separately for correctly specified subjects in **Exogenous** and **Endogenous** across periods. Panel A presents the first-best optimal action as a benchmark. The blue dashed line in Panel A represents the average first-best optimal action for correctly specified subjects in **Exogenous**. The red dashed line in Panel A represents the average first-best optimal action for correctly specified subjects in **Endogenous**. Panel B presents the average action a myopically optimizing Bayesian agent would take in the last period of the experiment as a benchmark. The simulations are conducted using each subject's prior beliefs about their abilities. The blue dashed line in Panel B represents the average simulated Bayesian action for correctly specified subjects in **Exogenous** whereas the red dashed line in Panel B represents the average simulated Bayesian action for correctly specified subjects in **Endogenous**.

Table A1: Estimation of Displacement Relative to the Simulated Bayesian Actions - Overconfident Subjects

	Panel A: Exogenous				Panel B: Endogenous			
	Dependent Variable: Δ_{BAYES}				Dependent Variable: Δ_{BAYES}			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
α	11.95^{*}	5.642^{*}	3.458	0.418	28.08^{***}	6.722^{*}	7.145^{*}	10.81^{**}
	(5.249)	(2.190)	(1.676)	(2.886)	(6.341)	(2.576)	(2.556)	(2.995)
Observations	20	20	20	20	22	22	22	22
Period	1	501	701	901	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for overconfident agents in **Exogenous** and **Endogenous**. Each column conducts the estimation $\Delta_{BAYES} = \alpha + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action.

Table A2: Estimation of Displacement Relative to the Simulated Bayesian Actions - Correctly Specified Subjects

	Panel A: Exogenous				Panel B: Endogenous			
	Dependent Variable: Δ_{BAYES}				Dependent Variable: Δ_{BAYES}			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
α	5.972	-2.152	0.590	-1.401	-5.526	5.099^{*}	2.381	2.634
	(4.842)	(2.978)	(3.318)	(2.364)	(4.736)	(2.415)	(2.060)	(1.700)
Observations	44	44	44	44	42	42	42	42
Period	1	501	701	901	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for correctly specified agents in **Exogenous** and **Endogenous**. Each column conducts the estimation $\Delta_{BAYES} = \alpha + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action.

Appendix B Experiment Instructions and Understanding Quizzes

INSTRUCTIONS

Welcome

You are about to participate in a decision-making experiment. In this experiment, you can earn a considerable amount of money, which will be paid to you through Venmo at the end of the experiment. The amount of money you earn will depend on your decisions. Therefore, it is in your best interest that you read these instructions carefully. If you need assistance, please raise your hand through the Zoom app. The experimenter will answer your question in a private chat.

The experiment consists of four parts. One of these parts will be randomly selected for payment at the end of the experiment. In the part that is randomly selected for payment, you can make either $25 or $0. In addition to your earnings from the experiment, you will receive a show-up fee of $10 for participating in the experiment. This means that at the end of the experiment you will receive either a payment of $35 (if in the randomly selected part you made $25) or $10 (if in the randomly selected part you made $0).

For each part of the experiment, you will be precisely instructed about your task.

Please put away your cell phone and do not interact with other participants throughout the experiment.

Instructions for Part 1

You will go through an IQ test in this part of the experiment. Tests similar to this are frequently used to measure intelligence.

The test consists of 10 questions, and you have 10 minutes to solve them. You should solve as many of the 10 questions as possible. Your earnings from this part of the experiment will be either $25 or $0. At the end of the experiment, we will randomly select one of your answers to the IQ test. If the selected answer is correct, you will earn $25 from this part of the experiment. This means the higher the number of correct answers, the more likely you will make $25 in this part of the experiment.

Instructions for Part 2

We conducted the exact same IQ test with other participants who previously, exactly like you, participated in an experiment at UCSB Experimental and Behavioral Economics Laboratory. We randomly selected 19 of these participants. Together with these 19 participants, you now form a group of 20 participants.

We constructed a ranking of this group based on the IQ test scores. The group member that scored the highest on the IQ test obtained rank 1. The group member with the second-highest score obtained rank 2, and so on. The group member with the worst performance on the IQ test got rank 20. In case of a draw between group members, the computer randomly decided who received the higher rank.

The computer then assigned you a color based on your ranking in your group. The top scoring members with ranking 1 to 4 are assigned Dark Green, the members with ranking 5 to 8 are assigned Light Green, the members with ranking 9 to 12 are assigned yellow, the members with ranking 13 to 16 are assigned light red, and the bottom scoring members with ranking 17 to 20 are assigned dark red.

How do you think you ranked on the IQ test?

In this part of the experiment, we are interested in how you think you ranked based on your IQ test score within your group of 20.

Your task is to submit your belief about how likely it is that you are assigned the color dark green, light green, yellow, light red or dark red based on the IQ test score rankings.

To indicate your beliefs, you will use a slider. Where you move the slider will represent your best assessment of the likelihood (expressed as a chance out of 100) that you are assigned one of these colors.

We will now go through an Explanation Stage to understand how the sliders work.

In this Explanation Stage, I would like you to enter a hypothetical subject's beliefs into the system.

Let's call this subject Ash.

Ash believes that their ranking is somewhere from 1st to 8th. However, they think it is more likely that they ranked from 1st to 4th rather than 5th to 8th. So, Ash believes they are more likely to be assigned the color dark green rather than light green.

Suppose, specifically, that Ash believes they are assigned dark green with a likelihood of 60 out of 100, and they are assigned light green with a likelihood of 40 out of 100.

Let us now enter Ash's beliefs into the system using sliders.

Now I would like you to do another example on your own.

In this exercise, you enter another hypothetical subject's beliefs into the system. Let's call this subject Blake.

Blake believes their likelihood of ranking from 13th to 16th is 40 out of 100, and ranking from 17th to 20th is 60 out of 100.

Please move the sliders to indicate Blake's beliefs and finalize.

Your Payment

You will be paid based on the accuracy of your belief. Your earnings from this part of the experiment will be either 25 or 0 USD, depending on how accurate your belief is about your color assignment based on the IQ test scores. This means the higher the likelihood your belief assigns to your actual color, the more likely you will receive $25.

If you understand this, you can click directly "Next". If you want to know the details of how we calculate your payments, please click "Details".

[Details on Your Payment]
After you state your belief, the computer will randomly draw a number k. This number is between 0 and 20,000. (More precisely, this number is drawn from a discrete uniform distribution on the interval from 0 to 20,000.) You will then receive $25 if the sum S is smaller or equal to k where S is the sum of the following elements:

- The squared deviation between the likelihood (out of 100) that you allocated to your actual color and 100 points.
- For *each* possible color that is not your actual color: The squared deviation between 0 points and the number of points that you allocated to that color.

The exact formula that we use to determine S is

$$S = \sum_{c \in C} (I(YourColor = c) \times 100 - L_c)^2$$

Where $I(YourRanking = c)$ is an indicator function that takes the value 1 if your color is c, $C = \{DarkGreen, LightGreen, Yellow, LightRed, DarkRed\}$ is the set of colors and L_i is the likelihood (out of 100) that you assign to color c.

While this formula might look complicated, the basic idea is very simple: you can secure the largest chance of winning $25 by reporting your most accurate belief about your assigned color.

Instructions for Part 3

Welcome to the main part of the experiment!

Introduction

At the beginning of this part of the experiment, we will assign a project to you. Your job is to act as a project manager for a company. Specifically, we will ask you to repeatedly recommend investment decisions to the company on your assigned project to maximize total profits from this project over multiple periods. The higher the profit you generate from this project, the more likely you will earn $25 from this part of the experiment.

Information on Projects

Projects have different qualities. Higher quality projects generate more profits.

Some projects are of higher quality than others. The project qualities can be any whole number between 0 and 100. The lowest possible project quality is 0, and the highest is 100. The higher the quality of your assigned project, the higher your profits are from the project.

You cannot choose or change your assigned project.

Although a high-quality project increases your profits, you cannot choose your project or its quality. We will randomly assign a project to you at the beginning of the experiment. You will be working on the same project that we assigned you until the end of the experiment.

You will not know your assigned project's quality.

The qualities of the projects vary between 0 and 100. All you will know about your assigned project's quality is that it can be any whole number between 0 and 100. You will not know your assigned project's quality until the end of the experiment.

Recap:

- You act as a project manager for a company in this part of the experiment
- We randomly assign you a project at the beginning of this part of the experiment
- You do not get to choose the project or its quality
- You work on the same project until the end of the experiment
- You repeatedly recommend to the company how much to invest into your assigned project

- Your assigned project's quality can be any whole number between 0 to 100, each number equally likely
- You do not know your project's quality until the end of the experiment
- Your goal is to maximize profits from the project you are assigned to

How do you maximize the profit from your assigned project in each period?

To make things easy for you, we designed the experiment so that it is straightforward to maximize the profit from your assigned project. You maximize your profit in each period by **recommending an investment amount that exactly matches what you think your project's quality is**.

Example.

Suppose you believe that your project's quality is 50. Then, you should recommend an investment amount of 50 to maximize the profit from the project in that period. Similarly, suppose you believe your project's quality is 74. In that case, you should recommend an investment amount of 74 to maximize the profit from the project in that period.

You can make sense of this simple profit-maximizing rule in the following way. If you have a high-quality project, you are better off investing a lot into that project as the return on that project is high. On the other hand, if you have a low-quality project, you are better off not investing too much into the project as the return on that project is low. Therefore, the profit-maximizing strategy is to ***match your recommended investment amount with the quality of the project***.

Now we will go through the details of calculating your profit when you recommend an investment decision. The instructions we will go through in the following pages might seem complex. However, please remember that we will <u>NOT</u> ask you to solve complex equations to maximize your profit during the experiment. The reason we provide these details is to ensure that you have a complete understanding of the experiment's rules.

The idea behind profit maximization is straightforward: <u>**recommend the investment amount that matches what you believe your project's quality is.**</u> You do not need to worry about maximizing your profit as long as you match your investment amount to what you think your project's quality is.

Please feel free to ask any questions along the way.

Details of Profit Calculations

The way we calculate profit in each period is straightforward. First, we will calculate the income you generate from the project and subtract the investment cost to calculate the profit. Then, we add a bonus of 5000 to ensure that no one ends up with a negative profit in the experiment.

$$Profit = ProjectIncome - InvestmentCost + 5000$$

As you see, profits have two main components: project income and investment cost. We will now go through each of these components individually.

Step 1: Project Income

1. ***Project Quality*** refers to the intrinsic quality of the project:
    - It will be a whole number between 0 and 100 in the experiment.
    - Higher quality projects generate higher incomes
2. ***Investment Amount*** refers to the amount you recommend the company to invest into the project:
    - You can choose any number between 0 and 100 as your Investment Amount
    - The higher the amount that is invested into your project, the higher the project income you generate
3. ***Your IQ Rank Score*** refers to your ranking in the IQ test you have completed in the previous part of the experiment. The Blue Table below describes how each ranking corresponds to a score:
    - The higher your ranking in the IQ test you have completed at the beginning of the experiment, the higher the project income you generate

Your Ranking	1-4	5-8	9-12	13-16	17-20
Your IQRANKSCORE	100	80	60	40	20

Specifically, we calculate the project income using the following equation:

$$ProjectIncome = ProjectQuality \times (InvestmentAmount + YourIQRankScore)$$

Step 2: Investment Cost

Investments you recommend to be made into the project have costs. The higher the amount you recommend to be invested, the higher the investment cost.

Specifically, we calculate the investment cost using the following equation:

$$InvestmentCost = \frac{(InvestmentAmount)^2}{2}$$

We can rewrite the full profit equation as

$$Profit = ProjectQuality \times (InvestmentAmount + YourIQRankScore) - \frac{(InvestmentAmount)^2}{2} + 5000$$

The green part of the profit equation is the income from the project and the red part of the profit equation is the investment cost.

If you have taken calculus, you can verify that the profit is maximized when you match the investment amount to the project's quality

$$InvestmentAmount = ProjectQuality$$

Recap:

- You maximize your profit in each period by matching your recommended investment amount to what you believe the project's quality is
- Higher quality projects generate higher profits
- We calculate an IQ rank score for you based on your ranking on the IQ test
- Higher IQ rank score generates higher profits

Will the company follow your investment recommendations immediately?

The company originally planned to invest an amount of 100 in each period on your project before your assignment. However, the company will immediately implement your recommended investment decisions and choose the amount you recommend in each period rather than the originally planned investment amount of 100.

Will you know how much profit you make after each investment decision?

A crucial point in the experiment is that you <u>will not</u> know how much profit you make after each investment decision. Instead, you will get an evaluation from your company if your profit is above or below your company's profit expectation. Since the company immediately implements your recommendations in each period during the experiment, the evaluations you get from the company will be based on your investment recommendations, not based on the company's originally planned investment amount of 100.

Recap:

- The company immediately implements your recommended investment decisions
- You will not know how much profit you make after each investment decision
- ...but you will know if you beat your company's profit expectation or not
- The evaluation you get from the company during the experiment is based on your recommended amount

How does the company determine its profit expectation?

The lowest possible profit you can generate in the experiment is 0, and the highest is 20,000. In each period, the company will randomly choose a profit amount, call it X, from the lowest possible profit amount (0) to the highest one (20,000) to expect from your project. If your profit is <u>at or above X</u> in a period, <u>you beat</u> your company's profit expectation. If your profit is <u>below X</u>, you <u>do not meet</u> your company's profit expectation.

Note that the higher your profit, the more likely you beat your company's profit expectation.

Will you know your company's profit expectations while making your decisions?

You will not know your company's profit expectation X before or after making your investment decision. The only information we will provide is if the profit you generate is above or below this profit expectation X.

What happens when you beat your company's profit expectation?

Once you make your last decision in the experiment, we will randomly select a period. If the profit based on your recommended investment decision beats the company's profit expectation in the randomly selected period, you earn $25 from this part of the experiment!

Recap:

- Your company chooses a number X between 0 and 20,000 as its profit expectation in each period, you will not know what X is
- The higher your profit, the more likely you beat your company's profit expectation
- If you beat your company's profit expectation in a randomly selected period, you earn $25 from this part of the experiment

We have established that you maximize profit in a period by matching your recommended investment amount with your project's quality. However, you do not know what your project's quality is! We will now go through how you can make some sophisticated guesses about your project's quality.

How can you make sophisticated guesses about your project's quality?

You can use your company's profit feedback to help you better understand your IQ rank score and your project's quality. Remember that your profits increase with your IQ rank score and your project's quality. <u>Hence any feedback that tells you that you beat the company's profit expectation is good news for your IQ rank score and your project's quality</u>.

<u>On the other hand, any feedback that tells you that you did not beat the company's expectations is bad news for your IQ rank score and your project's quality.</u>

To help you interpret the feedback that you get from the company, we will provide you with an expert statistician. In each period, the statistician will prepare a report for you, which you can use to make sophisticated guesses about your project's quality.

The Statistician's Report

These reports are going to look like the one on this page:

The Statistician's Report	
Your IQ Rank Score	**Project Quality**
20	75
40	71
60	69
80	66
100	65

The report is very straightforward to read. The statistician tells you:

- If your IQ rank score is 20, you should expect your project's quality to be 75.
- If your IQ rank score is 40, you should expect your project's quality to be 71, and so on.

Depending on what you believe your IQ rank score is, you may then make a sophisticated guess about your project's quality.

The statistician will update the report in each period incorporating the evaluations you receive from your company up until that period.

Instructions for Part 4

Please remember that at the beginning of the experiment, we assigned each participant in this session to a group with 19 other randomly selected people who had previously taken the same IQ test at UCSB Experimental and Behavioral Economics Laboratory.

We then constructed a ranking of each group based on the IQ test scores, and the computer assigned you a color based on your ranking in your group of 20.

The top scoring members with ranking 1 to 4 are assigned Dark Green, the members with ranking 5 to 8 are assigned Light Green, the members with ranking 9 to 12 are assigned yellow, the members with ranking 13 to 16 are assigned light red, and the bottom scoring members with ranking 17 to 20 are assigned dark red.

In this part of the experiment, we are again interested in how you think you ranked based on your IQ test score within your group of 20.

Your task is to submit your belief about how likely it is that you are assigned the color dark green, light green, yellow, light red or dark red based on the IQ test score rankings.

To indicate your beliefs, you’ll use a slider exactly as before.

Your Payment

You will be paid based on the accuracy of your belief. Your earnings from this part of the experiment will be either 25 or 0 USD, depending on how accurate your belief is about your color assignment based on the IQ test scores. This means the higher the likelihood your belief assigns to your actual color, the more likely you will receive $25.

If you understand this, you can click directly “Next”. If you want to know the details of how we calculate your payments, please click “Details”.

[Details are identical to Part 2’s Payment Details]

Understanding Quiz

Question 1.

Do you get to choose your project or its quality in the experiment?

- Yes, I choose both the project and its quality
- No, I do not get to choose either the project or its quality
- I only choose the project, but I don't get to choose its quality
- I do not get to choose the project, but I choose its quality

Question 2.

What type of decisions do you make on the project you are assigned to?

- I repeatedly give recommendations on how much the company should invest in the project
- I repeatedly give recommendations on how many projects the company should undertake
- I repeatedly give recommendations on whom to delegate the project
- I repeatedly give recommendations on which company should be responsible for the project

Question 3.

When do you learn your assigned project's quality?

- At the beginning of the experiment
- After my first investment decision
- Before my last investment decision
- At the end of the experiment

Understanding Quiz II

Question 1.

How do you maximize your profit in a period in the experiment?

- I match my recommended investment amount to the project's quality
- I match my recommended investment amount to my IQ Rank score
- I match my recommended investment amount to investment cost
- None of the above

Question 2.

Suppose you believe your project's quality is 62 in a period.

What investment amount maximizes your profits in that period?

- 31
- 62
- 93
- Not enough information to answer this question

Question 3.

Which of the below factors increase profits? [Multiple choice available.]

- Project's Quality
- My IQ Rank Score
- Investment Cost

Question 4.

What is your IQ rank score if you rank 1st in your group on the IQ test you have previously taken?

- 0
- 20
- 60
- 100

Question 5.

What is your IQ rank score if you rank 5th in your group on the IQ test you have previously taken?

- 0
- 20
- 60
- 100

Understanding Quiz III

Question 1.

When will the company implement your recommended investment decisions for each period?

- Immediately after I make my decisions
- Once I make my last decision
- At the beginning of the experiment, before I make any decision
- After I make my first decision, but before my last decision

Question 2.

What will we tell you after each investment decision you make?

- How much profit I make
- An evaluation from the company if my profit is higher than the company's profit expectation or not
- My assigned project's quality
- My IQ rank score

Question 3.

Before you make your last decision in the experiment, the evaluations you get from the company are based on which investment decisions?

- My recommended investment decisions
- The company's originally planned investment amount of 100
- Neither my recommended investment decisions nor the company's originally planned investment amount of 100
- Both my recommended investment decisions and the company's originally planned investment amount of 100

Understanding Quiz IV

Question 1.

How does the company choose its profit expectation in each period?

- It randomly chooses a number between the lowest and highest possible profit amounts (0 and 20,000)
- It uses historical data
- It uses investment costs
- It uses project's quality

Question 2.

How do we decide to pay you $25 in this part of the experiment?

- If my recommended investment decision generates a profit that beats my company's profit expectation in the first period
- If my recommended investment decision generates a profit that beats my company's profit expectation in the last period
- If my recommended investment decision generates a profit that beats my company's profit expectation in a randomly selected period
- None of the above

Understanding Quiz V

Question 1.

If you beat your company's profit expectations in a period, this is good news for

- Only your IQ rank score
- Only the project's quality
- Both your IQ rank score and the project's quality
- Neither your IQ score nor the project's quality

Question 2.

The Statistician's Report	
Your IQ Rank Score	**Project Quality**
20	75
40	71
60	69
80	66
100	65

If you think your IQ rank score is 60, what should you expect your project's quality to be according to the statistician's report?

- 75
- 71
- 69
- 66

Question 3.

If you think your IQ rank score is 80, what should you expect your project's quality to be according to the statistician's report?

- 75
- 71
- 69
- 66