Adjustment Dynamics During a Strategic Estimation Task*

Mel Win Khaw
Microsoft Corporation†

Luminita Stevens
University of Maryland

Michael Woodford
Columbia University

April 13, 2026

Abstract

We use data from a controlled laboratory experiment to test how decision-makers form beliefs in a strategic setting with an occasionally changing fundamental value. Payoffs depend on the exogenous fundamental as well as the group's average forecast. A range of models of expectation formation encompassing stochastic level-k and adaptive learning suggest limited strategic behavior and inattentive, experience-based forecasting. We find that 36% of subjects are best described as non-strategic level-0 forecasters, 44% as level-1 with limited strategic sophistication, and 10% as rational expectations forecasters whose sluggish adjustment reflects inattentiveness rather than shallow reasoning.

*Contact: melwinkhaw@gmail.com, stevens7@umd.edu, michael.woodford@columbia.edu.
†Listed affiliations are current ones; Dr. Khaw's work on this project took place while he was a postdoctoral research affiliate of Columbia University.

1 Introduction

Economic payoffs often depend on individual actions relative to the actions of others. How do people make choices in such interdependent settings? One possibility is that they employ a model-based, deductive approach: They make decisions given some understanding of the structure of their environment and based on some assumptions about the actions of others. Indeed, this is the quintessential approach to modeling decision-making in economics, and the resulting *rational expectations equilibrium* (REE) is the fundamental organizing principle of modern macroeconomic thought. However, empirical studies often find only limited evidence that decision-makers (DMs) engage in very sophisticated strategic thinking. For example, in a range of one-shot experiments testing strategic thinking in even the simplest environments, the typical finding is that while most subjects exhibit some degree of strategic thinking, the degree of strategic sophistication is quite limited, with most strategic subjects performing only one or two iterations when trying to infer what others might be thinking. Although this literature has found that experience with the task improves the degree of sophistication to some extent, the estimated levels are far from the infinite level of higher order thinking required by the REE model.[1]

Alternatively, decisions may be guided by past experience of actions and payoffs. People may pay attention to patterns in the data and gradually learn the statistical properties of their environment. Even if their choices resemble those that might be generated by a model-based approach, DMs might not explicitly consider what is driving these patterns, namely what model might generates the data they observe, or what these data imply about the actions of others. This has been the focus of a large learning literature in macroeconomics.[2]

Given enough depth of reasoning or accumulation of data, both the deductive and the experiential approach can bring an economic system to a (REE), and in some contexts, it may not matter *how* the system might get there—whether through gradual adaptation or model-based reasoning—especially if one believes that the system is already in a REE. However, this distinction becomes important in the presence of structural or policy changes that trigger transitions to new equilibria. For example, the welfare benefits of a policy change often depend on the speed with which the new equilibrium is reached. Such regime changes also make it possible for us to identify agents' reasoning modes. In the face of a credible and well-understood policy announcement, for instance, model-based agents will adjust immediately to the new regime, while experience-based agents will respond only

[1] Nagel (1995); Stahl & Wilson (1994, 1995). For a recent comprehensive review of this literature, see Mauersberger & Nagel (2018).

[2] See Evans & Honkapohja (1999) for the classic survey and more recently, Evans & Honkapohja (2013); Eusepi & Preston (2018).

gradually, as they accumulate data under the new regime.

In this article, we present a controlled laboratory experiment that gives participants the opportunity to use either type of approach, and we ask to what extent the subjects' behavior is indicative of an introspective, model-based way of reasoning, versus an experience-based approach. Our experimental design includes occasional regime changes over the course of many trials, that are clearly displayed on the participants' screen. This design allows for a sharp identification of the two modes of decision-making.

The experiment is a probability estimation task in a group setting in which the probability to be forecast depends on both an exogenous term and the average forecast of the group. The experimental design and the instructions provided to participants ensure that both modes of reasoning are useful to inform choices and also easy to apply, but neither is explicitly favored by the way we present the task. The participants make estimates for many periods and are shown realizations and payoffs in real time, so there is ample opportunity to extrapolate patterns. On the other hand, they are also given the information needed to form a strategic best response: they are shown the structure of the payoffs and the value of the fundamental, which they can use for strategic analysis. The experiment is designed to minimize adjustment costs, habits or other constraints on behavior, and it excludes any idiosyncratic shocks to individual payoffs, so that in principle it has none of the features that have typically been used to rationalize deviations from REE patterns of adjustment.

We identify subjective modes of reasoning by studying the extent to which individual forecasts conform to the predictions of a level-k model of rationality generalized to allow for various assumptions of how participants use the information on the screen and whether they attempt to forecast the forecasts of others. Our framework features three deviations from typical models of bounded level-k reasoning: *(i)* we incorporate adaptive expectations forecasting as the non-strategic behavioral rule, *(ii)* we allow for inattentiveness to the fundamental regime shifts in the strategic behavioral rule, and *(iii)* we explicitly model stochasticity in forecasts using entropy-based cognitive costs. Each of these features has been studied in isolation and found to be relevant in macro models. Here, we study their interaction in a controlled laboratory environment.

Both level-k models and adaptive expectations weaken the REE assumptions about people's abilities to reason or to learn about their environment, by lowering either the level of strategic sophistication or the ability to accumulate data in the learning process. As shown in applied work, these frameworks also open gaps between the outcomes of the thus-constrained economy and those of the REE economy. As a result, they have gained attention in macroeconomics, as a way to dampen equilibrium forces and forward looking behavior. See, for instance, Evans & Honkapohja (1999), Evans & Honkapohja (2013), Woodford (2003) Farhi

& Werning (2019), García-Schmidt & Woodford (2019), Vimercati, Eichenbaum & Guerreiro (2021). But which is a more relevant behavioral model for typical macroeconomic settings? And what speed of adaptation or degree of strategic sophistication should be assumed in those settings? Thus far, much of the work measuring the degree of strategic sophistication comes from *static* beauty contest games. Conversely, the work on adaptive expectations in dynamics settings has largely abstracted from strategic considerations in the learning process.

Our experimental results show that model-based reasoning fails even when the mapping is transparent. Forecasts are far from the REE, and they are noisy and biased, despite the relative simplicity of the task and of computing the REE. More importantly, approximately 36% of the subjects appear to be level-0. This finding contrasts results from a wide range of largely one-shot games, in which most subjects exhibit some degree of strategic thinking, with only a small fraction of participants (usually around 5%) behaving entirely non-strategic. The reason for this gap is that our experiment offers participants a useful alternative to strategic thinking: they can monitor the realizations for patterns, and indeed it seems that is what many of them do. Since depth of reasoning is not observable, forecasts that track patterns in the data may be mistakenly labeled as being strategic. For instance, if we estimate the level-k models with a naive level-0 rule that averages out to the mean of the expected outcome, the degree of estimated strategic reasoning shifts up by about 1.3 to 2.3 levels, relative to the level-k models with adaptive learning. Furthermore, we estimate that an additional 43% of participants are level-1. Overall, we uncover a prevalence of experience-based rather than model-based forecasting. Together with the limited degree of strategic sophistication, noisy adjustment leads to very slow aggregate convergence to the theoretical equilibrium. This suggests that macro models may need to incorporate much more bounded higher order beliefs than is currently the norm.

Our paper contributes to a large literature using controlled experiments to test rationality in strategic contexts. It provides further evidence that limitations in perception, information processing, and reasoning have significant and pervasive effects on economic decision-making. And it directly measures to what extent people's behavior incorporates attempts to predict the actions of others in their group. Laboratory settings are valuable for testing hypotheses regarding the formation of expectations since they allow the researcher to control not only the subjects' payoffs, but also the data generating process and the flow of information shown to the subjects. Hence, deviations from RE forecasts can be precisely measured both over the course of the experiment and in the cross-section. The literature has shown that arriving at a REE puts demands on individual decision-makers that may be unrealistic given the practical cognitive constraints that limit our ability to think strategically and to accumulate

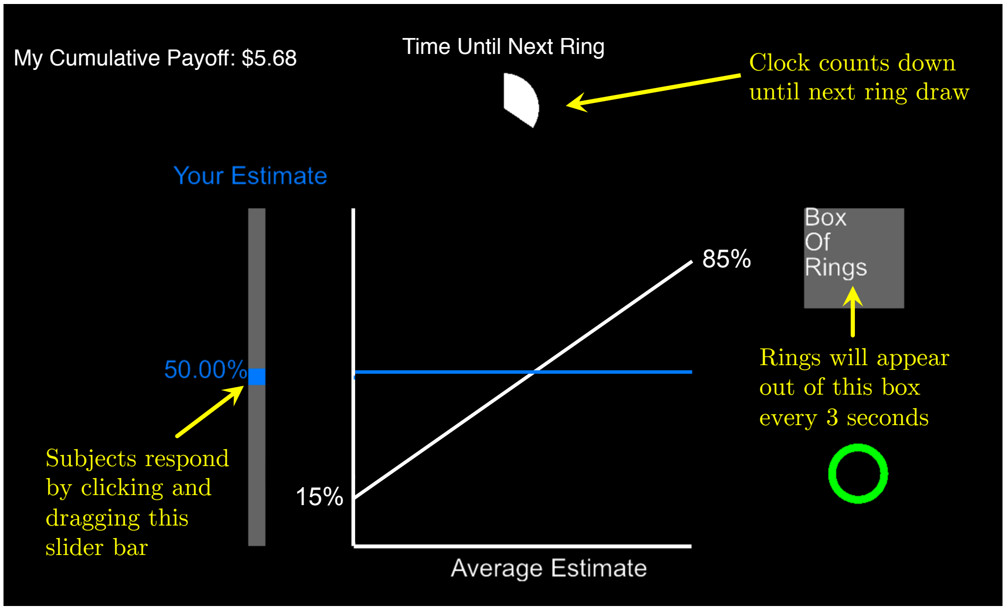

Figure 1: Screenshot of the experiment

data.

The experimental setting departs from the static environments that have been the focus of the beauty contest literature, bringing the design closer to the dynamic macro setup. We study a dynamic setting with a very long sequence of realizations of the exogenous random variable and with many repetitions of the task. Our experiment is designed to inform macroeconomic models more directly, hence the setting is akin to an economic decision-maker making the same kind of decision repeatedly, in a familiar context, much like a consumer purchasing a staple good or a firm updating the price of an existing product.

Section 2 describes the experiment and Section 3 discusses the main empirical patterns of adjustment. Section 4 presents the forecasting models we compare and Section 5 discusses the estimation results. Section 6 concludes.

2 Experiment

The experiment is an incentivized probability estimation task in which participants are placed in groups and are asked to repeatedly record their individual estimates of the probability of drawing a green ring from a virtual box that contains a mix of green and red rings.

2.1 Design

In each group, the true probability p_t of drawing a green ring on trial t depends on an exogenous term and on the group's average forecast on that trial:

$$p_t = z_t + \alpha \widehat{P}_t, \tag{1}$$

where z_t is the exogenous "fundamental," $\widehat{P}_t$ is the average of the group's current forecasts and $\alpha \neq 0$ measures the nature and strength of strategic interactions. The value of z_t remains fixed for a random number of trials and changes infrequently over time, as described below. The rational expectations equilibrium forecast, which emerges as a fixed point of equation (1), is given by

$$p_t^{REE} = \frac{z_t}{1-\alpha}. \tag{2}$$

The payoff per trial for each participant i is proportional to

$$v_{it} = 1 - (s_t - \widehat{p}_{it})^2, \tag{3}$$

where s_t is the realized state (equal to 1 for a green ring draw and 0 for a red one) and $\widehat{p}_{it}$ is the subject's recorded probability forecast for that trial. The total payment is a fixed participation amount plus an amount proportional to the total points earned in the session, $\sum_{t=1}^{T} v_{it}$.

Figure 1 shows a screenshot of what the participants see on their monitors. On the right side of the screen is the virtual box from which rings appear, one at a time, according to the hidden probability of green rings, p_t. In the center, a graph shows the individual best response probability on the y-axis, as a function of the group average belief $\widehat{P}_t$ on the x-axis. The line shown corresponds to the current fundamental z_t. The endpoints of the line (representing the best response probabilities associated with beliefs $\widehat{P}_t = 0$ and $\widehat{P}_t = 1$) are explicitly labeled. Whenever z_t changes, the line shifts and the extreme values are updated. This makes the strategic interaction explicit: participants are directly shown how their optimal forecast depends on what they think others believe for each possible value of the fundamental. This is explained to the subjects in the instructions preceding the experiment.

The left side of the screen has a vertical slider that participants use to choose a probability. Participants can freely move the slider up and down between 0 and 1, and a numeric display of their current forecast is also shown alongside the slider. Crucially, participants can map their chosen slider position onto the best-response graph in the center: for a given belief about the forecasts of others, participants can read off the slider what their implied best response would be. This feature helps make deductive reasoning as salient as possible. Hence, any limitations to the use of the model-based approach that we might uncover are likely to understate the limitations that would arise in practice, when people face much more complex environments.

Lastly, at the top of the screen, participants see their current cumulative payoff and the time remaining until the next ring is drawn from the box and displayed. For each trial, participants have three seconds to record their forecast. At the end of the three seconds, the computer draws a ring in accordance with equation (1), where the average forecast is computed based on the participants' slider positions in place at that time. Imposing the three-second time limit reflects practical considerations. Since the proportion of green rings in the box depends on the group average, having a time limit ensures that the experiment is not held up by any individuals taking a long time to record forecasts. This time limit also ensures that the task does not take too much time to complete. Each member of a group is shown the same information, but members cannot communicate with each other and must form their estimates on their own.

After each ring draw, the updated cumulative payoff is displayed, and there is a small probability δ of a change in z_t, in which case a new value is drawn independently from a uniform distribution that can take on three values: low, medium, and high. Whenever the value of z_t changes, the best-response line instantly shifts up or down to the new intercept, and the extreme values are updated.

2.2 Timeline

For clarity, the timeline of each trial and for each group is as follows:

1. A value of the fundamental z_t is realized and the associated best response function is displayed on the screen.

2. Each player in the group has three seconds to report their forecasts by adjusting the slider bar displayed on their screen.

3. After three seconds, a virtual box is formed with the proportion of green rings given by $z_t + \alpha\widehat{P}_t$, where the forecasts that make up the average forecast $\widehat{P}_t$ are the slider positions recorded by all the players in the group at the end of the three-second interval.

4. A ring is drawn from this box and displayed on the screen. Players' scores are updated as a function of the realized ring draw relative to each player's individual forecast, and the individual score is updated on the screen of each player.

5. If the maximum number of trials has been reached, the session ends. Otherwise, with probability δ, a new value of the intercept z is drawn and a new best response function is displayed on the screen; with probability $1 - \delta$, z and the best response function remain unchanged; the session continues with step 2.

2.3 Settings

We ran this experiment in the spring and summer of 2017 at Columbia University's Experimental Laboratory for the Social Sciences with undergraduate and master students. At the start of each session, subjects received written instructions and completed a 15-minute practice session. The participants were given all the information summarized in model (1) and instructions regarding the payoff structure and the incentives of the other participants. They completed the task at individual computer stations, and they were not allowed to communicate in any way with others in the group. They also wore earplugs to cover the sound of other participants' mouse clicks.

We conducted 18 sessions with 70 participants for a strategic complements treatment with $\alpha = 0.7$ and $z \in \{0.05, 0.15, 0.25\}$ and 15 sessions with 57 participants for a strategic substitutes treatment with $\alpha = -0.7$ and $z \in \{0.75, 0.85, 0.95\}$. These values ensure that regardless of an individual's beliefs about the forecasts of others, the individual's best response lies in the interval $[0.05, 0.95]$. Each group had between two and six participants. In this paper, we analyze the data from the strategic complements treatment.

2.4 Design Discussion

The experimental design's key distinguishing feature is that it leaves participants free to choose how to approach the forecasting problem, while making the underlying strategic structure fully transparent. Since participants observe both the sequence of ring realizations and the best response function that summarizes the game, we can observe which approach they adopt and measure their performance without potentially underestimating how well they *could* do the task.

The design thus bridges two experimental classes: On the one hand, the setup belongs to the class of experiments of Gallistel, Krishan, Liu, Miller & Latham (2014), Ricci & Gallistel (2017), and Khaw, Stevens & Woodford (2017), which display the sequence of ring realizations, enabling the participants to observe and learn from the patterns in the data. On the other hand, we also incorporate a strategic component, make public the exogenous fundamental, and display the best response function. These elements connect our design to strategic expectation-formation experiments in the tradition of Nagel (1995) and Fehr & Tyran (2008), where outcomes depend endogenously on others' beliefs.

The experiment is also unusual in that it generates very long time series punctuated by occasional regime shifts: Each session consists of $T = 1,000$ trials with infrequent shifts in the fundamental ($\delta = 0.5\%$). This gives the participants ample experience with the task while also giving us a long enough time series of observations, over the course of which we can observe both how expectations change in response to a change in the exogenous intercept, and how they are subsequently revised in response to experience in the new regime.

Overall, the structure of the task allows for rich learning and coordination dynamics. Because the true probability depends on the group's average forecast, participants face strategic uncertainty about others' beliefs. While shifts in the fundamental are made salient through corresponding shifts in the best-response function, participants do not directly observe the group's average belief and must infer it from realized outcomes over time. The feedback from the aggregation of beliefs to the outcomes allows for the possibility of fast coordination toward the rational expectations equilibrium, but also of persistent dispersion, delayed adjustment, or self-reinforcing deviations, depending on how and whether participants form and update beliefs about others.

3 Patterns of Adjustment

This section describes how participants update their estimates over time and when there is a change in the fundamental.

3.1 Incomplete Adjustment

How do subjects update their estimates in response to infrequent regime changes? Individual forecasts aggregate to a sluggish, hump-shaped, and very incomplete response function. Consider a change in the fundamental at the beginning of some trial t. This change decays at rate δ, which is the probability that a new fundamental value is drawn (independently) on each subsequent trial. Since the change is displayed on the participants' screen immediately, the REE forecast jumps fully, on impact, per equation (2), and then decays at the same rate as the fundamental.

Figure 2a compares the responses of our participants to the REE benchmark by plotting the expected change in the average forecast at time $t + h$ in response to a change in the fundamental at time t, relative to the expected change in the absence of an innovation to the fundamental. The REE response is normalized to 1 on impact. The average response is only

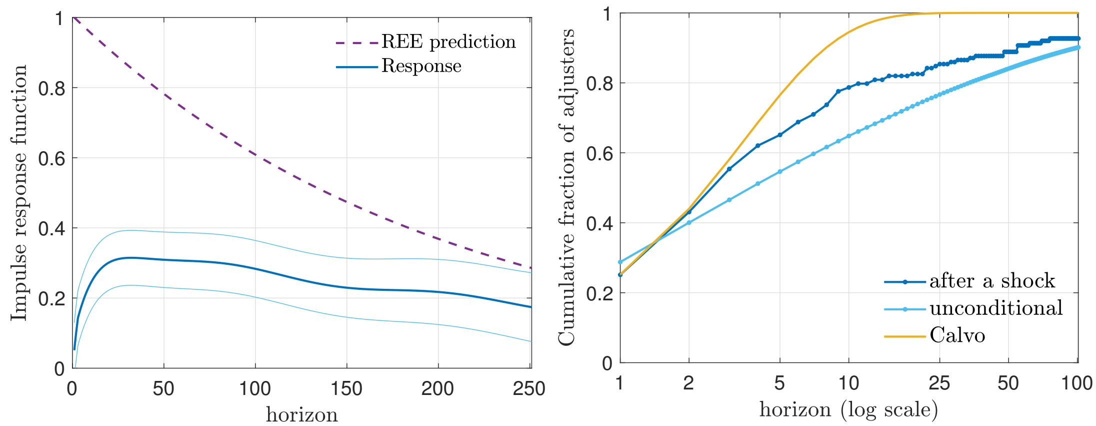

(a) Impulse response function

(b) Cumulative fraction of adjusters

Figure 2: Adjustment versus the rational expectations prediction

Notes: (a) The impulse response function for the average group forecast to a shift in the fundamental and its 95% confidence interval, compared to the REE response. (b) The cumulative fraction of adjusters at different horizons following a shift in the fundamental, and unconditionally. The fraction is computed as the number of subjects first adjusting within h periods divided by the total number of subjects, pooled across data from all sessions. The x-axis uses a log scale, to show early adjustment more clearly. The RE model predicts 100% adjustment on impact. The Calvo model is parameterized to match the average frequency of adjustment in the data.

a small fraction of the REE response, it peaks at approximately 31% roughly 30 periods after the change in the fundamental, and then decays very slowly. This evidence is consistent with an extensive literature in macroeconomics that has documented sluggishness in aggregate forecasts and actions (Coibion & Gorodnichenko, 2015; Carroll, Crawley, Slacalek, Tokuoka & White, 2020).

3.2 Delay

Driving the incomplete adjustment is the fact that participants do not notice or choose not to respond to the shift right away, and instead adjust in a staggered manner (*e.g.*, Calvo (1983); Mankiw & Reis (2002)). Indeed, only 25% of forecasts are adjusted on impact, even though the change is displayed on the subjects' screens instantaneously. Moreover, the cumulative rate at which participants adjust their forecasts following an innovation in the fundamental value is only moderately higher than the unconditional cumulative rate of adjustment, across all trials, as shown in Figure 2b. The figure plots the cumulative probability that subjects will have adjusted their forecast at least once within a given number of periods, both unconditionally and following a shift in the exogenous state. Compared with the adjustment of an exponential parameterized to match the mean frequency of adjustment

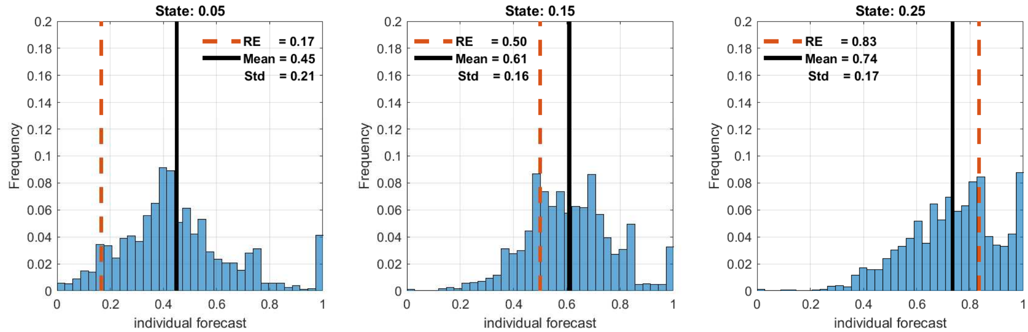

Figure 3: Distribution of individual forecasts around the mean forecast, versus the RE forecast, for each value of the fundamental state z_t. `Std` indicates the standard deviation of the individual forecasts for each state.

in our data, we see that participants adjust at a significantly lower rate in later trials.

3.3 Dispersion and Bias

The extensive margin of adjustment is only partially responsible for the slow overall adjustment. The empirical response function would adjust significantly faster than what we observe if participants adjusted to the optimal forecast when they did adjust. Instead, there is both dispersion and bias in forecasts. As a result, the aggregated response is much more sluggish than the frequency of adjustment alone would imply.

Figure 3 plots the distribution of individual forecasts for each value of the fundamental. Across all sessions, individual forecasts differ from the REE prediction by significant amounts in absolute value. Moreover, forecasts are biased: on average they are 8.7% higher than the REE forecast. The magnitude of the deviations falls toward the high end of the range of deviations typically reported in the experimental literature, such as the work that studies mis-pricing in experimental asset markets spurred by Smith, Suchanek & Williams (1988). The biggest gaps occur when the fundamental value is lowest: subjects tend to over-estimate the optimal forecast in the low state by 28 percentage points and under-estimate it in the high state by nearly 10 percentage points.

Figure 4 provides an alternative visualization of the deviations in the subjective forecasts. Given the individual forecasts and the realized exogenous state, we can infer what individual subjects believe the group average to be on each trial, under the assumption that they are attentive to the state and use equation (1) to form their forecasts. This implied belief is given by $E_t^i \widehat{P}_t = (\widehat{p}_t^i - z_t)/\alpha$. The figure plots the mean of these implied beliefs $(\widehat{P}_t - z_t)/\alpha$ against the actual mean forecast $\widehat{P}_t$, for all trials. When what people believe the mean to be

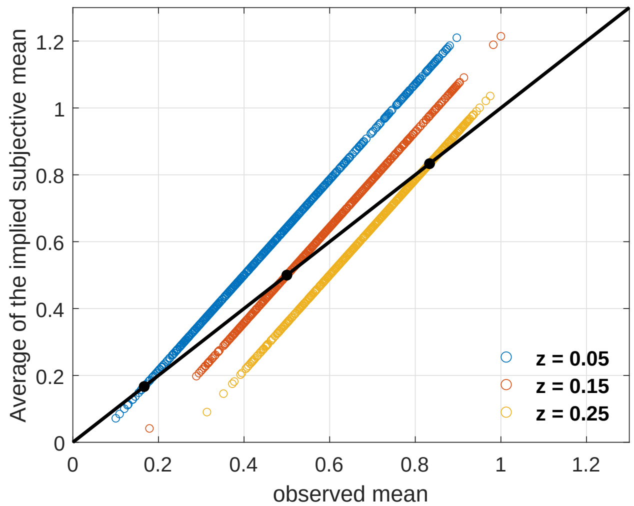

Figure 4: The implied average of subjective forecasts, $(\widehat{P}_t - z_t)/\alpha$ against the actual average $\widehat{P}_t$. The black dots mark the REE for each state.

coincides with the actual mean, we have a REE. The figure shows this is virtually never the case. Instead, interpreted through the lens of model (1), participants tend to either over-estimate or under-estimate the actual mean. Moreover, some beliefs cannot be rationalized with equation (1) at all, since doing so would generate beliefs that the average forecast is greater than 1. This motivates us to consider modeling beliefs in a generalized framework that allows for deviations from model (1) in terms of how and whether participants use the mapping between the fundamental and the equilibrium probability to form beliefs.

4 Modeling Adjustment

The experiment is designed to allow us to test two classes of belief formation: model-based and experience-based. First, participants are shown the best response function — the linear relationship between the average forecast and the individual's optimal estimate, for each fundamental. Using this information, they can compute their response to any individual belief about the group forecast, however they might form that belief. Moreover, the rational expectations equilibrium (REE) forecast itself does not require a complex calculation, so participants could quite easily use the rational group forecast, per equation (2). Second, even if participants do not use a model-based approach to predict the next ring realization—for example, because they do not understand the best response graph or its implications—they can still perform well, and, at least in principle, approach the REE forecast, by monitoring the patterns in the data: they observe a long series of draws, and they face a fundamental that changes only rarely, which gives them ample opportunity to learn. The sequence of realizations offers a valuable alternative, experience-based method of forecasting.

Let $\widehat{p}_{it}$ denote the forecast recorded by DM i on trial t. The DM's payoff per trial is proportional to $1 - (\widehat{p}_{it} - s_t)^2$, where s_t is equal to 1 with probability p_t and 0 otherwise. Let $\mathcal{M}^i$ denote a DM's subjective model of how p_t is determined. To identify the forecasting approach that subjects seem to prefer, we consider stochastic, dynamic generalizations of bounded-rationality deterministic level-k models as different possibilities for $\widehat{p}_{it}$ and $\mathcal{M}^i$. The model variants differ in terms of how decision-makers use the information available to them and in the extent to which they attempt to forecast the forecasts of others. We do not model how a DM might come to use one subjective model over another; instead, we assess which subjective model fits the data best.

4.1 Stochastic Level-k with Learning

Consider the typical specification of level-k reasoning in the experimental games literature, which allows for differing levels of strategic sophistication, or depths of reasoning. Level 0 involves no consideration given to what others might do. Level 1 is assumed to be constrained to doing only one round of reasoning: forecasters assume that others behaves as if they were level-0 decisionmakers. Each higher level is assumed to be able to do one additional round of reasoning. Hence, when modeling the forecasts of others, the most a level-k forecaster can do is to assume that others are doing at most k-1 iterations. As the depth of reasoning increases, beliefs approach the REE. This approach to relaxing REE has been used extensively to describe beliefs in static beauty contest games, as noted in the introduction.

In the canonical level-k model, it is assumed that the non-strategic level-0 decision-makers draw random forecasts on the support allowed, while higher level forecasters deterministically best-respond to the expected value of the lower level forecasts. However, our data do not exhibit such a dichotomy across subjects. Instead, all forecasts are noisy, but at the same time, they are less likely to be at the extremes than near the middle of the unit interval. This motivates us to consider a control cost specification to generate *endogenous* stochasticity of forecasts. We furthermore allow individuals to take into account the randomness in other people's forecasts: in the level-k hierarchy, each level generates a noisy best response to the noisy lower-level response, as in quantal response equilibria (QRE) (McKelvey & Palfrey, 1995, 1998).

The approach exploits the fact that noisy forecasts can be modeled using an information-theoretic framework. Suppose we have a candidate model $\mathcal{M}^i$ that implies a forecast m_{it} in period t. If the DM faced no other cognitive or control costs associated with recording the forecast implied by this subjective model, they would set $\widehat{p}_{it} = m_{it}$. Instead, with additional cognitive costs, the DM draws forecasts from a *distribution* of potential forecasts. If the DM exerts no effort, their recorded forecast will be a random draw from the allowable support, regardless of the model forecast. This would generate forecasts that reduce to the standard "naive" level-0 assumption. By exerting effort, individuals can concentrate their forecasts and record a forecast that improves their expected payoff. But doing so is costly. The DM's objective, excluding control costs, is to report forecasts to maximize the expected cumulative payoff

$$\mathcal{V}_i = \sum_{t=0}^{T} E_{it}\left[1 - (\widehat{p}_{it} - p_t)^2\right], \tag{4}$$

where we omit discounting and terms independent of the subjective forecast. We assume that the control cost is proportional to the KL divergence from the uniform distribution on the unit interval, which simplifies to

$$\mathcal{C}_i(\pi) \;=\; \lambda_i \int_0^1 \pi(x) \ln \pi(x)\, dx, \tag{5}$$

where λ_i denotes the unit control cost. Hence, the optimization problem becomes choosing the sequence of distributions $\{\pi_{it}\}$ to maximize the expected payoff net of the costs associated with drawing forecasts from these distributions. The problem reduces to a sequence of static optimization problems: In each period, the DM draws their forecast from a distribution that maximizes their expected payoff net of control costs:

$$\max_{\pi \in \Delta} \int_0^1 \left[1 - (\widehat{p} - m)^2\right] \pi(\widehat{p}|m) d\hat{p} - \mathcal{C}_i(\pi), \tag{6}$$

where Δ is the set of probability measures on the unit interval.

For each m, the solution to (6) is

$$\pi(\widehat{p} \mid m) = \frac{\exp\left\{-\frac{(\widehat{p}-m)^2}{\lambda}\right\}}{\int_0^1 \exp\left\{-\frac{(\tilde{p}-m)^2}{\lambda}\right\} d\tilde{p}} \cdot \mathbf{1}_{\{\widehat{p}\in[0,1]\}}. \tag{7}$$

This is a truncated normal distribution on $[0, 1]$ with pre-truncation mean m and variance $\lambda/2$. We denote this distribution by $\mathcal{N}^T(m, \lambda/2; 0, 1)$. The solution parallels the logit form familiar from the quantal response equilibrium literature, with the control cost parameter governing the degree of noise in choices. As the control cost approaches zero, forecasts concentrate around the target m; as it grows large, the distribution approaches the uniform default. This specification nests as a limiting case the canonical level-k model, in which level-0 forecasters have infinite control costs while higher level forecasters have zero control costs. Here, we estimate control costs for all players.

Although expected payoffs under quadratic losses depend only on the mean of the forecast distribution, stochasticity plays a non-trivial role in our setting. Because forecasts are constrained to lie in the unit interval, noise interacts with truncation, so that the mean of the choice distribution need not coincide with the deterministic target m. Thus, DMs take stochasticity into account not because they value higher moments of the forecast distribution per se, but because noise influences the implied mean forecasts and their dynamics.

Level 0 In the canonical level-k model, the non-strategic level-0 DMs draw random forecasts independently across trials. However, in our experiment, many participants do not seem to approach the forecasting session as a sequence of independent static games. Instead, we observe propagation of beliefs across trials within a session. Forecasts tend to be serially correlated and also seem to correlate with recently realized ring draws. To take this into account, we model the dynamic propagation of beliefs by allowing for adaptive, experience-based learning.

Suppose that non-strategic DMs watch the ring realizations and apply a constant-gain learning algorithm to form their forecast. These DMs do not use information about the fundamental or about the forecasts of others to form their forecast, but they do watch the rings as an *alternative* to model-based deductive forecasting. The recorded forecast of the level-0 DM is a draw from the distribution

$$\widehat{p}_{it}^{(0)} \sim \mathcal{N}^T \left(m_{it}^{(0)}, \lambda_i/2; 0, 1 \right), \tag{8}$$

$$m_{it}^{(0)} = (1 - \gamma_i)\; m_{i,t-1}^{(0)} + \gamma_i \left(s_{t-1} + b_i \right), \tag{9}$$

where $m_{it}^{(0)}$ denotes the deterministic constant gain forecast, the initial belief is $m_{i1}^{(0)} = 0.5 + b_i$, γ_i is the individual's learning parameter, s_{t-1} is the realized ring draw, which they use in their learning, and b_i is an individual-specific bias. The model forecast becomes a weighted average of the history of rings (shifted by b_i) and the initial belief.

The bias b_i captures persistent deviations from the unbiased benchmark. When $b_i = 0$, the initial belief equals 0.5 and, in the long run, $m_{it}^{(0)}$ tracks the weighted history of the ring realizations. When $b_i \neq 0$, the forecast is persistently shifted away from the unbiased benchmark. The learning parameter γ_i governs responsiveness to recent ring realizations, independently of the bias: a subject can learn quickly from the data (γ_i large) while at the same time maintaining a persistent bias. In the special case where $\gamma_i = 0$ and $b_i = 0$, the model reduces to $m_{it}^{(0)} = 0.5$ in all periods, which corresponds to the standard naive level-0 specification in which forecasts fluctuate around the ex-ante expected value of the probability.

Level 1 The level-1 forecasters understand the structure of the environment but think that everyone else forms beliefs according to the level-0 model. They stochastically best respond to their beliefs about the mean of the stochastic level-0 forecast. They are subject to cognitive frictions with unit cost λ_i, and they understand that level-0 forecasts are also subject to such costs. Each level-k forecaster uses their own λ_i when simulating the behavior of lower levels. This mirrors the standard QRE assumption that people use their own parameters when simulating the behavior of others.

In what follows, let $\widehat{P}_{it}^{k-1|k}$ denote the level-k belief of forecaster i about the mean of the level-$(k-1)$ forecasts. For a level-1 forecaster, this is the expected value of the level-0 forecast distribution:

$$\widehat{P}_{it}^{0|1} = E\left[\hat{p}_{it}^{(0)}\right], \tag{10}$$

where the expectation is taken over the distribution given by (8), but with forecaster i's gain, bias, and control cost parameters. To simulate level-0 forecasts, the level-1 DM monitors the sequence of ring realizations and applies the constant gain learning rule in (9) (with their own parameters γ_i and b_i)

The resulting level-1 forecasts are governed by

$$\widehat{p}_{it}^{(1)} \sim \mathcal{N}^T\left(m_{it}^{(1)}, \lambda_i/2; 0, 1\right) \tag{11}$$

$$m_{it}^{(1)} = z_t + \alpha \widehat{P}_{it}^{0|1}. \tag{12}$$

Higher Levels Higher level forecasts are defined recursively, with each level generating noisy best responses to the average of the lower level's stochastic forecasts. Specifically, each higher-level forecaster constructs estimates of expected forecasts $\widehat{P}_{it}^{0|k}, \widehat{P}_{it}^{1|k}, ..., \widehat{P}_{it}^{k-1|k}$, and sets

$$\widehat{p}_{it}^{(k)} \sim \mathcal{N}^T\left(m_{it}^{(k)}, \lambda_i/2; 0, 1\right) \tag{13}$$

$$m_{it}^{(k)} = z_t + \alpha \widehat{P}_{it}^{k-1|k}. \tag{14}$$

Within the model, the form of updating used by non-strategic forecasters is common knowledge, just as the payoff structure and informational environment are common knowledge in standard level-k models.

Finally, we consider forecasters whose beliefs correspond to a fixed point of the level-k

hierarchy: beliefs about others' forecasts are self-consistent, and agents take into account others' expected randomization in their strategic reasoning. We refer to these as QRE forecasters, since the term "rational expectations" conventionally implies both infinite strategic depth and no other cognitive costs that might give rise to stochasticity in forecasts.

In the estimation, we allow for heterogeneity across individuals in all behavioral parameters. Each participant i is characterized by a discrete depth of reasoning k_i, a unit control cost λ_i, a learning gain parameter γ_i, and belief bias b_i. For each participant, all parameters are jointly estimated and the Bayesian Information Criterion (BIC) is used to select among competing specifications, trading off fit and parsimony.

The different model components generate distinct implications for forecasting behavior. Variation in the depth of reasoning k_i affects how strongly individual forecasts respond to the average forecasts of others and to changes in the best-response mapping, with higher levels exhibiting responses closer to the rational expectations benchmark. Variation in the learning gain γ_i governs the speed of belief revision in response to new ring realizations, with higher values generating faster adjustment following shifts in the underlying probability, while $\gamma_i = 0$ yields forecasts that display no systematic dependence on past realizations. The bias b_i shifts the level of forecasts away from the unbiased benchmark, independently of learning speed. Finally, the control cost λ_i governs the dispersion of forecasts around the implied target, with higher values generating noisier forecasts. These differences allow the data to discipline both the informational assumptions and the degree of strategic sophistication underlying individual forecasts.

4.2 Incorporating Inattentiveness

So far, persistence in forecasts arises solely from experience-based learning at the level-0 stage, which anchors the forecasts of all higher levels. This may lead to an over-estimate of the incidence of non-strategic behavior if persistence instead reflects sluggish adjustment by more sophisticated forecasters. To address this, we extend the model to allow for persistence in higher-level forecasts, independently of level-0 learning or strategic depth, by introducing inattentiveness to the fundamental value z_t. The key idea is that individuals may not continuously monitor the fundamental, even if they understand its role in determining the best-response function. As a result, forecasts may adjust sluggishly following changes in z_t, independently of learning about the behavior of others. This type of friction has been used in the monetary literature to generate sluggish nominal adjustment (Mankiw & Reis, 2002) and here we integrate it with level-k reasoning.

We model inattentiveness to the fundamental by allowing strategic DMs to update information about z_t infrequently, while using knowledge of the fundamental's data generating process to update its evolution between observations. We assume all DMs begin fully informed about the initial fundamental value. On each subsequent trial, individual i fails to update information about the fundamental with probability $\phi_i \in (0, 1)$, and with probability $1 - \phi_i$ the individual observes the current realization z_t. Hence, the perceived fundamental of a DM who last observed the fundamental h periods ago is

$$\widehat{z}_{t|t-h} = (1-\delta)^h z_{t-h} + \left(1 - (1-\delta)^h\right) \bar{z}, \tag{15}$$

where $\bar{z} = 0.5$ is the mean of the redraw distribution for the fundamental.

The subjective model forecast of an inattentive level-k decision maker with information lag h is

$$m_{iht}^{(k)} = \widehat{z}_{t|t-h} + \alpha \widehat{P}_{iht}^{k-1|k} \tag{16}$$

and the recorded forecast is a draw from

$$\widehat{p}_{iht}^{(k)} \sim \mathcal{N}^T \left(m_{iht}^{(k)}, \lambda_i/2;\, 0, 1 \right). \tag{17}$$

The term $\widehat{P}_{iht}^{k-1|k}$ is the expected value of the simulated level-$(k-1)$ forecast distribution, which now depends on the DM's own information lag h.

The role of inattentiveness differs across depths of reasoning: The forecasts of the level-0 DMs are not affected by inattentiveness to the fundamental, since these forecasters rely solely on the sequence of ring realizations. Hence, the level-0 expected forecast remains as described earlier. For a level-1 decision maker, inattentiveness affects beliefs about the fundamental z_t but not beliefs about others' forecasts, since this DM believes all others are level-0 forecasters. Hence, for level-1 DMs, the forecast of others continues to be the mean of the distribution given by (8), independent of the DM's own information lag.

For higher levels of reasoning, inattentiveness affects both beliefs about the fundamental and beliefs about the forecasts of others. DMs correctly anticipate the distribution of information vintages when making forecasts. From the perspective of a DM with information lag h, the fundamental at a more recent date $l < h$ (used by the cohorts with more recent information) is

$$E\left[z_{t-l} \,|\, z_{t-h}\right] = (1-\delta)^{h-l} z_{t-h} + \left(1 - (1-\delta)^{h-l}\right) \bar{z}. \tag{18}$$

The forecaster with lag l then further discounts this toward $\bar{z}$ over the l periods until t, so that the belief of a DM with lag h about the cohort-l perceived fundamental on trial t is

$$E\left[\hat{z}_{t|t-l} \,|\, z_{t-h}\right] = (1-\delta)^h z_{t-h} + \left(1 - (1-\delta)^h\right) \bar{z} = \hat{z}_{t|t-h} \;\text{ for } l < h. \tag{19}$$

That is, from the perspective of a DM with lag h, all cohorts with fresher information ($l < h$) have the same expected perceived fundamental, equal to the DM's own perceived fundamental. The same expression holds for $l = h$, where both DMs observed z_{t-h} directly.

For the cohorts who observed the fundamental at an earlier date, the DM with information lag h has no way of knowing what they saw when, so this DM assumes that they have converged to the mean fundamental, $\bar{z}$. The key insight to gain tractabiliy involves making a forward-rational but backward-uninformed assumption for handling computational complexity when DMs reason about others with different information lags: forecasters can project forward from their observations but use ergodic priors for DMs with more stale information. Hence,

$$E\left[\widehat{z}_{t|t-l} \mid z_{t-h}\right] = \begin{cases} \widehat{z}_{t|t-h}, & \text{if } l \leq h, \\ \bar{z}, & \text{if } l > h. \end{cases} \tag{20}$$

Consider a level-2 DM with information lag h. This DM considers the different cohorts of lower level forecasters, indexed by their information lag $l \in \{0, 1, 2, \ldots\}$. Given the perceived fundamentals for each cohort, the level-2 DM with information lag h can construct the model forecast of level-1 DMs for each cohort l based on

$$m_{ilt}^{(1)} = E\left[\widehat{z}_{t|t-l} \mid z_{t-h}\right] + \alpha \widehat{P}_{it}^{0|1}. \tag{21}$$

The level-2 DM understands that level-1 forecasts in cohort l are drawn from $\mathcal{N}^T(m_{ilt}^{(1)}, \lambda_i/2;\, 0, 1)$, and that cohort l has mass $(1-\phi_i)\phi_i^l$ in the population. Hence, the distribution of level-1 forecasts is a mixture of truncated normals:

$$\widehat{p}_t^{(1)} \sim \sum_{l=0}^{\infty} (1-\phi_i)\phi_i^l \cdot \mathcal{N}^T(m_{ilt}^{(1)}, \lambda_i/2). \tag{22}$$

The level-2 DM's belief about the mean level-1 forecast $\widehat{P}_{iht}^{1|2}$ is therefore the mean of this mixture. Using (20), this simplifies to a weighted average of two terms:

$$\widehat{P}_{iht}^{1|2} = \sum_{l=0}^{h} (1-\phi_i)\phi_i^l \cdot E\left[\mathcal{N}^T(\widehat{z}_{t|t-h} + \alpha \widehat{P}_{it}^{0|1}, \lambda_i/2)\right] + \sum_{l=h+1}^{\infty} (1-\phi_i)\phi_i^l \cdot E\left[\mathcal{N}^T(\bar{z} + \alpha \widehat{P}_{it}^{0|1}, \lambda_i/2)\right] \tag{23}$$

which simplifies to

$$\widehat{P}_{iht}^{1|2} = (1-\phi_i^{h+1}) \cdot E\left[\mathcal{N}^T(\widehat{z}_{t|t-h} + \alpha \widehat{P}_{it}^{0|1}, \lambda_i/2)\right] + \phi_i^{h+1} \cdot E\left[\mathcal{N}^T(\bar{z} + \alpha \widehat{P}_{it}^{0|1}, \lambda_i/2)\right]. \tag{24}$$

That is, the level-2 DM's belief is a weighted average of two truncated normal means: one based on their own perceived fundamental $\widehat{z}_{t|t-h}$, weighted by the probability $(1 - \phi_i^{h+1})$ that a random level-1 forecaster has information at least as fresh; and one based on the ergodic mean $\bar{z}$, weighted by the probability ϕ_i^{h+1} that a random level-1 forecaster has staler information.

For levels $k \geq 3$, an additional complication arises: level-$(k-1)$ DMs with different information lags hold different beliefs about level-$(k-2)$ forecasts, and a fully rational level-k DM would account for this heterogeneity. To maintain tractability, we assume that when simulating lower levels, DMs attribute their own current beliefs about even-lower-level forecasts to all forecasters at the intermediate level. This parallels our earlier assumption that DMs use their own λ_i and ϕ_i when simulating the behavior of others.

Under this assumption, the structure generalizes to arbitrary levels of reasoning. For a level-k DM with information lag h, beliefs about lower levels are constructed recursively. Let $\widehat{P}_{iht}^{j|k}$ denote the belief of a level-k DM with lag h about the mean level-j forecasts, for $j < k$. The belief about level-0 forecasts is independent of the DM's information lag:

$$\widehat{P}_{iht}^{0|k} = \widehat{P}_{it}^{0|1} = E\left[\widehat{p}_{it}^{(0)}\right]. \tag{25}$$

For $j \in \{1, ..., k-1\}$, the level-k DM recognizes that level-j forecasters best respond to their own beliefs about level-$(j-1)$ forecasts, using their own perceived fundamental. Applying the same logic as for level-2, the level-k DM's belief about the mean level-j forecast is

$$\widehat{P}_{iht}^{j|k} = (1-\phi_i^{h+1})\cdot E\left[\mathcal{N}^T\left(\widehat{z}_{t|t-h} + \alpha\widehat{P}_{iht}^{j-1|k},\ \lambda_i/2;\ 0,1\right)\right] + \phi_i^{h+1}\cdot E\left[\mathcal{N}^T\left(\bar{z} + \alpha\widehat{P}_{iht}^{j-1|k},\ \lambda_i/2;\ 0,1\right)\right]. \tag{26}$$

The level-k DM's own model forecast is then

$$m_{iht}^{(k)} = \widehat{z}_{t|t-h} + \alpha\widehat{P}_{iht}^{k-1|k}, \tag{27}$$

and the recorded forecast is drawn from $\mathcal{N}^T(m_{iht}^{(k)}, \lambda_i/2;\ 0, 1)$.

Finally, an inattentive QRE forecaster with information lag h has perceived fundamental $\hat{z}_{t|t-h}$ given by equation (15). Their forecast satisfies:

$$m_{h,t}^{QRE} = \hat{z}_{t|t-h} + \alpha \cdot E_h[\bar{P}_t], \tag{28}$$

where $E_h[\bar{P}_t]$ is their expectation of the average forecast. Using the model's belief structure, a forecaster with lag h believes that fraction $(1 - \phi^{h+1})$ of others have information at least as fresh as their own, and the fraction ϕ^{h+1} have staler information. Hence,

$$E_h[\bar{P}_t] = (1-\phi^{h+1})m_{h,t}^{QRE} + \phi^{h+1}\bar{m}^{QRE}, \tag{29}$$

where $\bar{m}^{QRE} = \bar{z}/(1-\alpha)$ is the QRE forecast for someone whose perceived fundamental has converged to $\bar{z}$.

Substituting and solving for the fixed point yields:

$$m_{h,t}^{QRE} = \frac{\hat{z}_{t|t-h} + \frac{\alpha}{1-\alpha}\phi^{h+1}\bar{z}}{1-\alpha(1-\phi^{h+1})}. \tag{30}$$

When $\phi = 0$, this reduces to the standard attentive REE forecast $m_t^{QRE} = z_t/(1-\alpha)$. For $\phi > 0$, the forecast is a weighted combination of the response to the perceived fundamental and the unconditional mean, with the weights depending on the information lag and the degree of strategic complementarity, α.

Inattentiveness is an additional source of sluggishness in the higher level forecasts themselves: if higher level forecasters do not immediately notice regime shifts, they will introduce delays in the average forecast. Moreover, since this inattentiveness is stochastic, the delays will be staggered, generating smooth sluggishness.

5 Estimation Results

In this section, we discuss the best-fitting estimates across models, and the importance of the generalized frameworks in identifying the degree of strategic sophistication.

We infer the population type frequencies from data fitting exercises at the subject level. For each subject and each model type, we estimate the best-fitting set of parameters to minimize the Bayes Information Criterion (BIC). The parameters to be estimated are the depth of strategic reasoning k_i, the control cost λ_i, the adaptive learning coefficient on ring-watching γ_i, the level bias b_i, and the inattentiveness probability ϕ_i (for inattentive higher level forecasters). We allow for heterogeneity in these parameters across subjects, and in the estimation, we seek to isolate the relevance of these different components in accounting for deviations of the subjective forecasts from the REE prediction.

5.1 Estimation Method

We estimate each model separately for each subject using maximum likelihood, and use the Bayesian Information Criterion (BIC) to select among models and levels of reasoning.

For a model with n parameters estimated on T trials,

$$\text{BIC} = -2\log\mathcal{L} + n\log(T), \tag{31}$$

where $\mathcal{L}$ is the maximized likelihood. The number of parameters n varies across specifications: naive attentive models estimate only the effort cost λ; adaptive attentive models estimate λ and the learning gain γ; and inattentive models additionally estimate the inattentiveness probability ϕ for levels $k \geq 1$. For the REE benchmark, only λ is estimated for the attentive versions, and λ and ϕ are estimated for the inattentive variants.

The likelihood for each trial is based on a truncated normal distribution on $[0, 1]$. Given a model forecast m_t and effort cost λ, the forecast is drawn from a normal distribution with mean m_t and variance $\lambda/2$, truncated to the unit interval. The log-likelihood contribution of observing forecast $\widehat{p}_t$ is

$$\log \ell_t = \log\varphi\left(\frac{\widehat{p}_t - m_t}{\sigma}\right) - \log\left[\Phi\left(\frac{1-m_t}{\sigma}\right) - \Phi\left(\frac{0-m_t}{\sigma}\right)\right], \tag{32}$$

where $\sigma = \sqrt{\lambda/2}$, and $\varphi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal pdf and cdf, respectively.

We select each subject's level of reasoning k by minimizing BIC within each model specification, and we determine the best-fitting model for each subject by comparing BIC across models.

We estimate parameters using constrained optimization with parameter bounds $\lambda \in [0, 1]$ for the effort cost; $\gamma \in [0, 1]$ for the learning gain; and $\phi \in [0, 0.99]$ for the inattentiveness probability, where the upper bound excludes the degenerate case in which the fundamental is never observed.

For the inattentive models, computing the likelihood requires summing over the distribution of information lags h, where a DM with lag h last observed the fundamental h periods ago. Because each DM's realized information lag is unobserved, the likelihood of forecast $\widehat{p}_t$ is a mixture over all possible lags h, weighting the truncated normal density evaluated at each lag-specific model forecast $m_t(h)$ by the probability $(1-\phi)\phi^h$ of that lag. We truncate this mixture at lag $\bar{h}$, treating all DMs who have not observed the fundamental for $\bar{h}$ or more periods as if their last observation was exactly $\bar{h}$ periods ago. To ensure consistent approximation quality across different values of ϕ, we set $\bar{h}$ adaptively such that $\phi^{\bar{h}+1} < 0.001$, *i.e.*, $\bar{h} = \left\lceil \frac{\log(0.001)}{\log(\phi)} \right\rceil - 1$, with a floor of 5 and ceiling of 200. This ensures that this truncation affects less than 0.1% of the probability mass for $\phi \lesssim 0.97$.

TABLE I: Levels of strategic reasoning: best fitting model

Level	Count	Fraction
L0	25	0.43
Attentive L1	18	0.36
Inattentive L1	19	0.27
Attentive L2	1	0.01
Inattentive L2	3	0.04
Inattentive L5, L6	2	0.03
Inattentive QRE	2	0.03
Total	70	1.00

Notes: Maximum likelihood estimates of the level of strategic reasoning k for each subject, under the best-fitting adaptive level-k model. L0 indicates non-strategic behavior ($k = 0$); QRE indicates quantal response equilibrium ($k \to \infty$).

TABLE II: Levels of strategic reasoning: attentive model

Level	Count	Fraction
L0	30	0.43
L1	31	0.44
L2	4	0.06
L3	1	0.01
L5	2	0.03
QRE	2	0.03
Total	70	1.00

Notes: Maximum likelihood estimates of the level of strategic reasoning k for each subject, under the attentive adaptive level-k model.

5.2 Results by Model Type

Table **??** collects the results for the models with attentive and inattentive forecasters.
YOU ARE HERE

5.2.1 The importance of experience-based learning

The most striking contrast is between naive and adaptive models in their implied distributions of strategic sophistication. Under the naive specifications, which anchor level-0 (L0) beliefs at the uninformative prior of 0.5, a meaningful fraction of subjects are classified as rational expectations (RE) forecasters (26% in Naive Attentive, and 20% in Naive

TABLE III: Levels of strategic reasoning: inattentive model

Level	Count	Fraction
L0	31	0.43
L1	31	0.44
L2	3	0.04
L3	1	0.01
L5	1	0.01
L5	1	0.01
QRE	2	0.03
Total	70	1.00

Notes: Maximum likelihood estimates of the level of strategic reasoning k for each subject, under the inattentive adaptive level-k model.

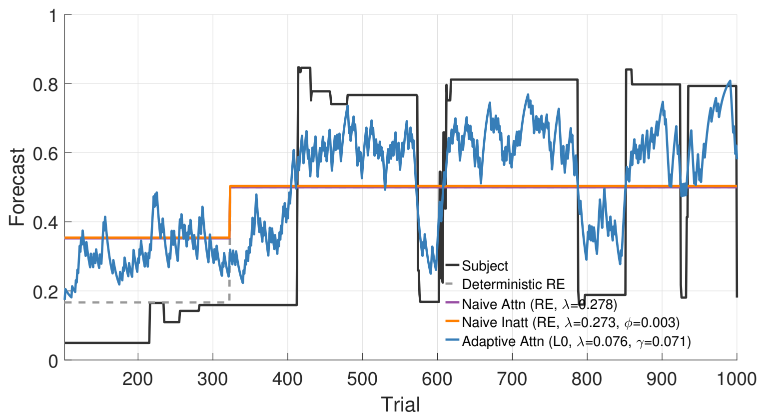

Figure 5: Adaptive learning versus rational expectations

Notes: Time series for a single subject classified as RE under the naive specifications but L0 under adaptive specifications.

Inattentive). But when we allow for the possibility that forecasters learn from the observed ring realizations, the level distribution shifts dramatically: under Adaptive Attentive, 53% of subjects are classified L0, 31% are L1, with 3% at RE. This shift suggests that adaptive learning can explain patterns that otherwise might be labeled as strategic reasoning. For many DMs, forecasts track the underlying fundamental quite well not due to sophisticated strategic reasoning and the use of a correct model of the data generating process, but rather through atheoretical inductive updating on the ring realizations.

To illustrate with a concrete example, Figure 5 shows the time series of subjective forecasts and model-implied forecasts for a subject who is classified as forming rational expectations under both naive specifications, but as level-0 under the adaptive specifications. The adaptive L0 model tracks the subject's forecasts much better, by updating beliefs based on

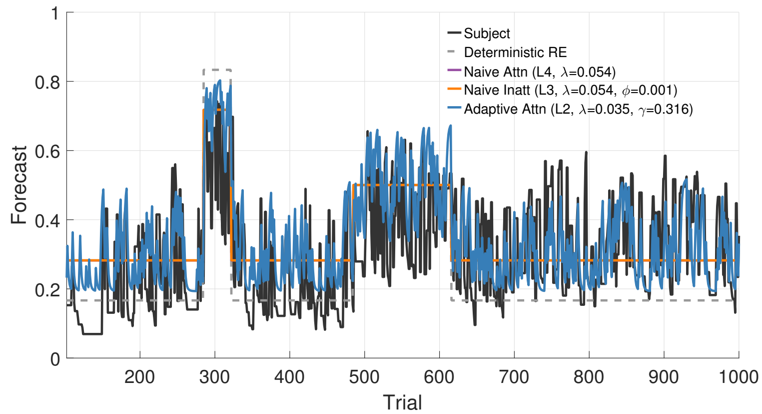

Figure 6: Adaptive learning versus strategic sophistication

Notes: Time series for a single subject classified as strategically sophisticated under naive attentive but L2 under adaptive attentive. What appears as deeper reasoning in naive models reflects adaptive learning.

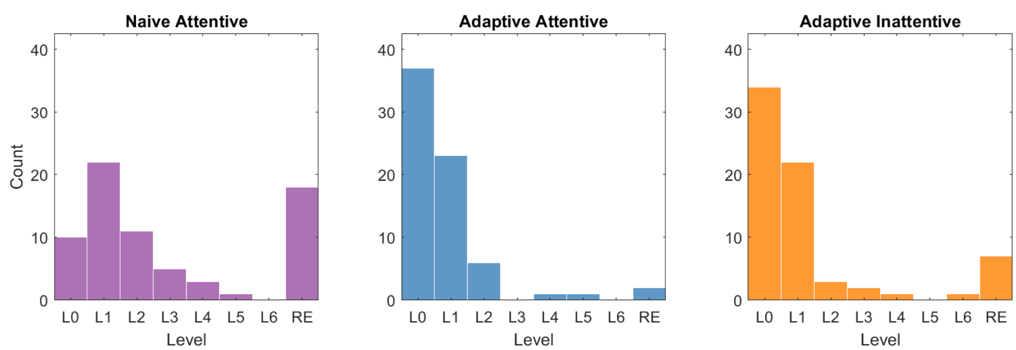

Figure 7: Distribution of estimated levels of strategic reasoning

Notes: Histograms show the number of subjects classified at each level under three model specifications.

observed ring realizations, capturing dynamics that the naive model can only rationalize by assuming the subject forms rational expectations about others' behavior.

While the gap is most striking between RE and L0, the pattern also arises at other levels of strategic sophistication: Figure 6 shows the forecasts of a subject who is categorized as L4 and L3 under the naive specifications, but in fact very closely tracks the L2 adaptive forecast. What appears to be higher reasoning in a naive model is in fact fast adaptive learning.

Overall, the distribution of levels of strategic sophistication shifts to the left when introducing learning, as shown in Figure 7.[3] Moreover, learning also lowers the estimated control costs. As shown in Figure 8, for most subjects, the estimated effort cost λ is smaller once

[3]The estimation caps the finite levels at L7 for all models, but raising this cap does not change the estimated levels.

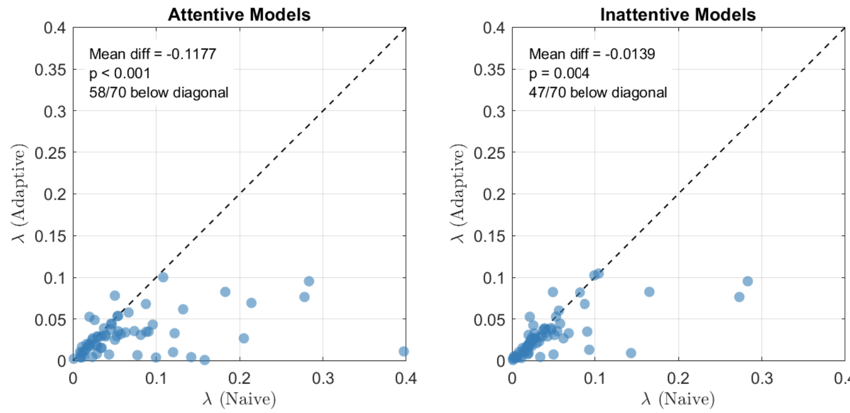

Figure 8: Learning gain versus effort in adaptive models

Notes: Each point represents one subject. Points below the diagonal indicate lower estimated effort cost λ under the adaptive specification.

one allows for learning, regardless of the presence or degree of inattentiveness. This suggests that without the learning channel, naive models absorb actual updating into the noise term.

5.2.2 The role of inattentiveness to the fundamental

Allowing for inattentiveness to the fundamental is also critical as it helps to correctly identify strategically sophisticated forecasting. To illustrate, Figure 9 shows the forecasts of a subject that is L0 in the attentive specifications, but L3 in the adaptive inattentive specification (which also fits the subjective forecasts the best by far). Without accounting for inattentiveness, this subject's sluggish adjustment to changes in the fundamental is attributed to shallow strategic reasoning.

Overall, the distribution of levels of strategic sophistication shifts to the right when introducing inattentiveness, as can be seen when comparing the last two panels of Figure 7. In particular, the inattentive specification identifies considerably more RE subjects than the attentive models.

Inattentiveness is also an important source of sluggishness that simultaneously allows subjects to respond strongly to the data they see. This has systematic effects on the estimated parameters. First, the distribution of learning gains shifts to the right, relative to the adaptive attentive estimation, as shown in Figure 10: Subjects who appeared to be slow to update under attentive specifications are now estimated to learn quickly from the information they observe, but to observe some of it infrequently.

Second, inattentiveness also changes the relationship between estimated parameters. In

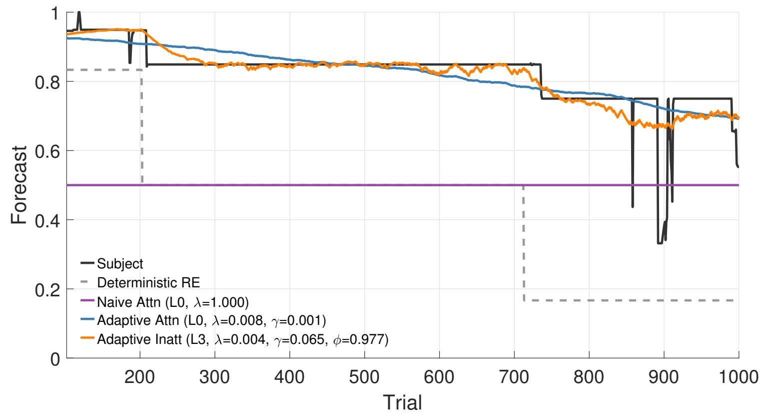

Figure 9: Inattentiveness helps identify strategic reasoning

Notes: Time series for a single subject whose sluggish adjustment is attributed to shallow reasoning (L0–L1) under attentive specifications but to inattentiveness with deeper reasoning (L3) in the Adaptive Inattentive model.

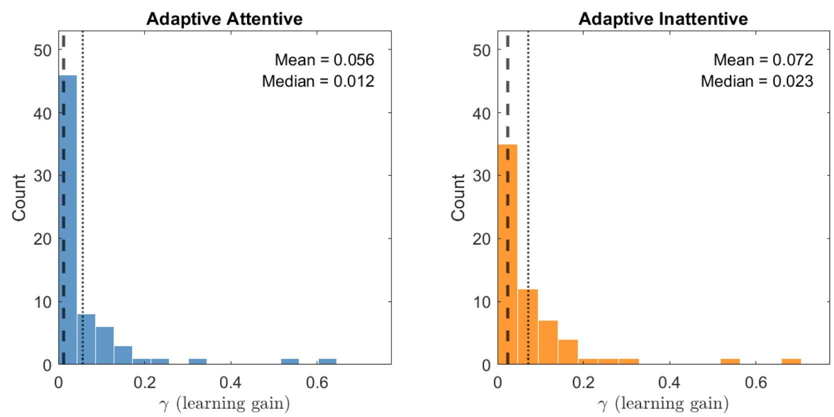

Figure 10: Distribution of estimated learning gain parameters

Notes: The dashed lines indicate medians; the dotted lines indicate means. Allowing for inattentiveness shifts the distribution of γ to the right.

the attentive specifications, high sensitivity to the data needs to be offset by large noise in forecasts to match the observed behavior, generating a positive correlation between γ and λ. This correlation disappears in the adaptive inattentive specification, as shown in the scatter plots in Figure 11, leaving the control cost and the learning speed as independent frictions.

Figure 12 shows that there is also no correlation between inattentiveness and learning or between inattentiveness and the effort cost. This suggests that inattentiveness is not simply a proxy for noisier or slower forecasts. Instead, these three parameters capture distinct channels: costly precision (λ), learning about level-0 behavior (γ), and attention to fundamentals (ϕ) all contribute independently to explaining forecast behavior.

Overall, introducing inattentiveness ($\phi > 0$) improves fit within both the naive and adap-

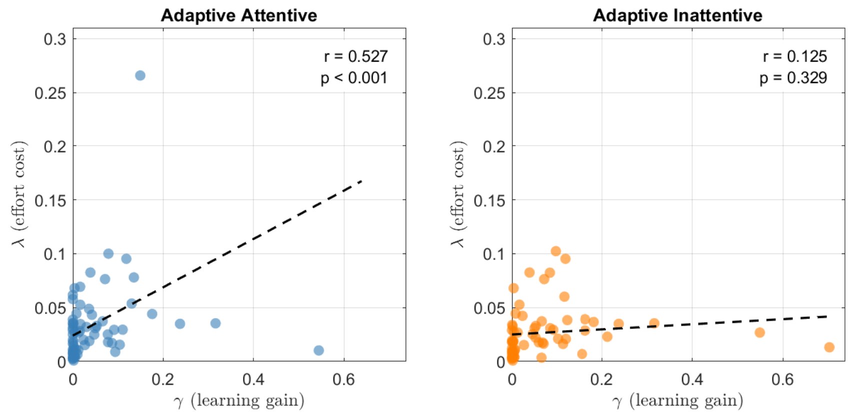

Figure 11: Learning gain versus effort with and without inattentiveness

Notes: Each point represents an individual subject. Panels also report the correlation r between the learning gain γ and the effort cost λ and the corresponding p-value.

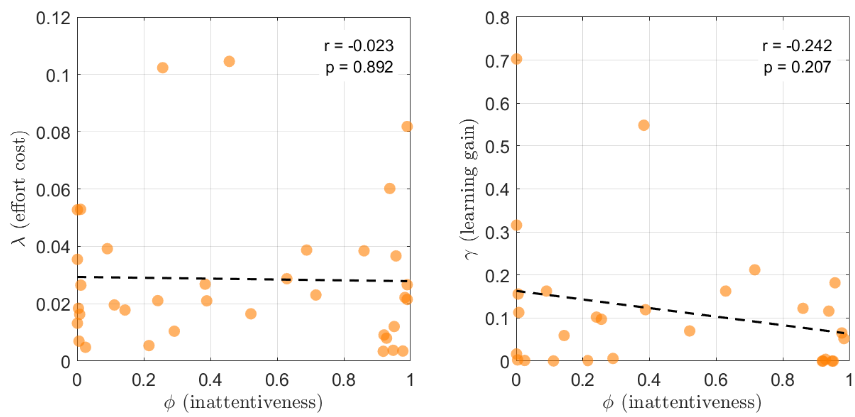

Figure 12: Parameter relationships in the Adaptive Inattentive model

Notes: Left panel: inattentiveness (ϕ) vs. effort cost (λ). Right panel: inattentiveness (ϕ) vs. learning gain (γ). Each point represents one subject. Correlation coefficients and p-values reported in upper right.

tive frameworks, as reflected in lower total BIC values for inattentive relative to attentive specifications. Moreover, the estimated inattention probabilities are substantial, and show considerable heterogeneity depending on model and level, as shown in Figure 13 which plots the distribution of estimated inattentive parameters in both the naive and the adaptive models. This suggests that imperfect observation of the fundamental is an empirically important feature of forecasting.

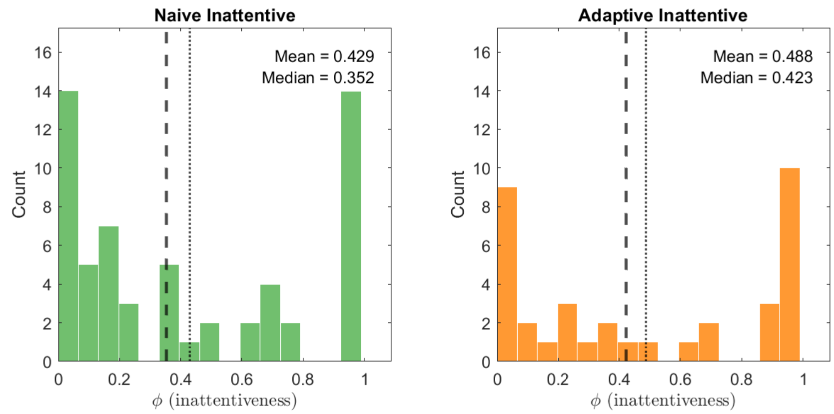

Figure 13: Distribution of inattentiveness in Naive and Adaptive models

Notes: Histograms show estimated ϕ for subjects with $k \geq 1$. The dashed lines indicate medians; the dotted lines indicate means.

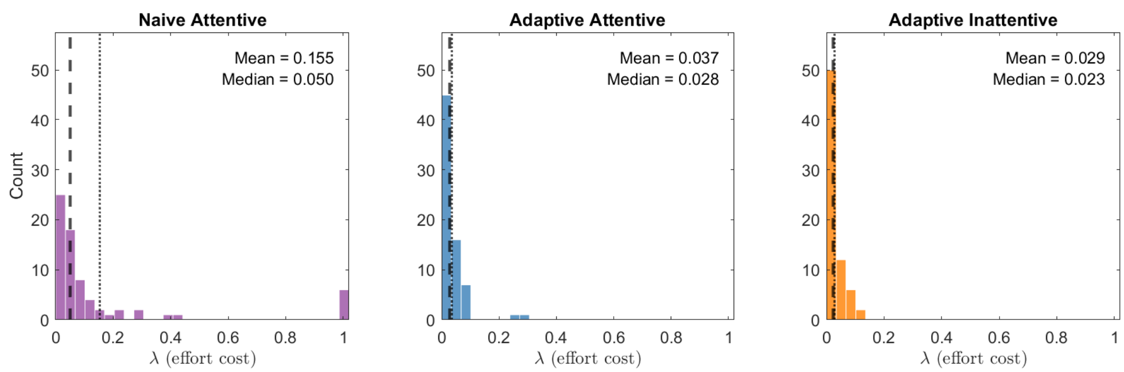

Figure 14: Distribution of effort costs across models

Notes: Dashed lines indicate medians; dotted lines indicate means. All panels use common bin widths for comparability.

5.2.3 The importance of control costs

Figure 14 displays the distribution of estimated effort costs λ_i across subjects for each model specification. Several features stand out. First, there is substantial heterogeneity: estimated λ_i ranges from near zero to above 0.15, indicating that some subjects implement their forecasts with high precision while others exhibit considerable noise. Second, the distributions are right-skewed, with most subjects clustered at low values but a tail of higher λ subjects in each specification. Third, the distributions are broadly similar across models, though allowing for learning significantly shrinks the estimated control costs.

Allowing for heterogeneous noise is not merely a matter of improving model fit. Modeling the endogeneity of noise via the control cost framework offers an interpretation for the noise we observe in the subjective forecasts. Rather than treating forecast errors as pure measurement error, λ gives them a behavioral interpretation: precision is costly, so subjects optimally choose imperfect implementation.

On the identification front, the control costs play an important role in absorbing idiosyncratic, trial-to-trial variation, thus allowing the other parameters to isolate the systematic, history-dependent patterns of forecasting. Introducing λ allows us to distinguish model failure from imperfect implementation: Without it, any deviation from the deterministic forecast would constitute model failure. With it, we can distinguish subjects who are strategically sophisticated but implement noisily from those whose behavior is inconsistent with the model altogether. We saw this was an important feature of allowing for inattentiveness, in order to correctly identify the RE subjects, and it is also important in the case of control costs.

Counterfactual To demonstrate these effects we conduct a counterfactual estimation in which we fix λ at the pooled median of our baseline estimation ($\bar{\lambda} = 0.023$) and then estimate the rest of the parameters as before. In this exercise, the subjects who are high-λ in the baseline and are now forced down to the median are a proxy for what would happen to noisy implementers if we did not model their imprecision.

Table IV reports the shifts in estimated levels of strategic reasoning when moving from the heterogeneous $\{\lambda_i\}$ to the fixed $\bar{\lambda}$ estimation and Table V reports how the changes in estimated parameters when λ is constrained correlated with the baseline levels of λ. We find that subjects with high imprecision are systematically reclassified as having lower depth of reasoning, and that learning and attention are also affected. Intuitively, since we are forcing noisy forecasters to appear precise, their dispersed forecasts are attributed to other channels: lower sophistication, more inattention, or noisier learning.

5.3 Best Fit Across Models

Table VI reports the best fitting models for the different subjects. First, 36% of subjects are *Level-0* forecasters. This fraction is far above the typical results from the game theoretic literature which tends to find approximately 5% of subjects as best described by non-strategic behavior. The gap reflects experience-based learning which is a valuable alternative to model-based reasoning. At the other extreme, none of our subjects are attentive RE forecasters, but 10% of them are *inattentive* RE forecasters. This underscores the importance of disentangling inattention from strategic sophistication: without inattentiveness, only 2 subjects would be labeled as RE and the others would all become Level-0 Adaptive forecasters. Finally, 44%

TABLE IV: Classification Changes Under Fixed $\bar{\lambda}$

	High-λ Subjects			Low-λ Subjects		
Model	Down	Same	Up	Down	Same	Up
Naive Attentive	21	34	0	0	15	0
Naive Inattentive	10	28	1	1	29	1
Adaptive Attentive	6	30	4	2	26	2
Adaptive Inattentive	2	31	2	4	30	1

Notes: High-λ subjects have baseline $\lambda > \bar{\lambda}$ (forced down to median); low-λ subjects have baseline $\lambda \leq \bar{\lambda}$ (forced up to median). Down/Up refer to changes in estimated level of reasoning towards lower/higher sophistication.

TABLE V: Correlations Between Baseline λ and Parameter Changes Under Fixed $\bar{\lambda}$

Model	Corr(ΔLevel, λ)	Corr($\Delta\gamma$, λ)	Corr($\Delta\phi$, λ)
Naive Attentive	-0.379 ($p = 0.001$)	–	–
Naive Inattentive	-0.640 ($p = 0.000$)	–	$+0.306$ ($p = 0.018$)
Adaptive Attentive	-0.073 ($p = 0.551$)	-0.779 ($p = 0.000$)	–
Adaptive Inattentive	-0.229 ($p = 0.056$)	-0.301 ($p = 0.018$)	$+0.042$ ($p = 0.816$)

Notes: Each cell reports the correlation between a subject's baseline λ and the change in the indicated parameter when λ is fixed at the sample median $\bar{\lambda}$. p-values in parentheses.

of subjects are either attentive or inattentive *adaptive* forecasters, 77% of which are level-1. Hence, we uncover a prevalence of experience-based forecasting, with very limited levels of strategic reasoning, despite the design of the experiment, which makes strategic reasoning not only salient but also transparent and simple.

The limitations in perception, information processing, and reasoning that we uncover have sizable consequences for participants' payoffs and hence welfare. Figure 15 shows a bar chart of the the losses suffered by participants, according to our best fitting models, relative to the Bayesian full information benchmark. Typical losses range from 20% to 25%, with a median loss of 24%. We decompose the losses into the model contribution and the contribution from control costs,[4] and find that control costs contribute $1-1.5\%$ to the losses.

These losses are an order of magnitude larger than the consumption-equivalent welfare costs of business cycles or of suboptimal policy typically found in sticky-price or sticky-

[4]Given the quadratic scoring rule, the expected payoff loss due to λ is approximately equal to the variance of the realized forecasts, $\sigma_i^2 = \lambda_i/2$. Due to truncation to the unit interval, the payoff losses are only approximately equal to the variance, but the truncation effects are minor.

TABLE VI: Subject classification by model-level category

Category	N	%	k values
Naive Level-0	3	4.3	0
Naive Attentive Level-k	3	4.3	1
Naive Inattentive Level-k	4	5.7	1
Adaptive Level-0	22	31.4	0
Adaptive Attentive Level-k	16	22.9	1, 2, 4, 6
Adaptive Inattentive Level-k	15	21.4	1,2,3
Inattentive RE	7	10.0	∞

Notes: Subjects classified by best-fitting model-level combination. Categories collapse distinctions that are irrelevant at certain levels: at Level 0, attention does not matter (subjects do not observe the fundamental); at RE, naive vs. adaptive does not matter (subjects do not use Level-0 forecasts).

information models, which range from 0.1 to 1 percent of steady-state consumption. While the laboratory scoring rule is not directly comparable to consumption equivalents, the magnitude of the losses underscores that the behavioral frictions we estimate, namely limited strategic sophistication, slow experiential learning, and inattentiveness to regime shifts, impose meaningful costs on individual decision-makers, even in a setting where the mapping from fundamentals to optimal forecasts is simple and transparent.

Figure 16 summarizes the distributions of estimated parameters across subjects. Most subjects exhibit low implementation costs (λ concentrated near zero), suggesting that imprecise execution of intentions is not a major source of forecast errors. Learning gains (γ) are similarly right-skewed, with the majority of adaptive subjects making only modest belief revisions in response to new information. In contrast, inattentiveness (ϕ) is dispersed across the full unit interval, with a notable mass near the upper bound, indicating that some subjects update their beliefs very rarely. The bottom row confirms that the three parameters are largely uncorrelated, consistent with each capturing a distinct behavioral channel: the cost of precise implementation, the speed of belief revision, and the frequency of information processing.

Finally, Figure 17 shows the theoretical impulse response functions for the different models, given the best-fitting estimated parameters for each subject. Under *Naive Attentive*, forecasts jump immediately to the level-k best response and thereafter track the decay of the fundamental at rate $(1-\delta)^h$. Under *Naive Inattentive*, the immediate response is geometrically smoothed: forecasts catch up to the attentive benchmark at rate $(1-\phi)$ per period. But because of the higher estimated strategic reasoning for these subjects, the re-

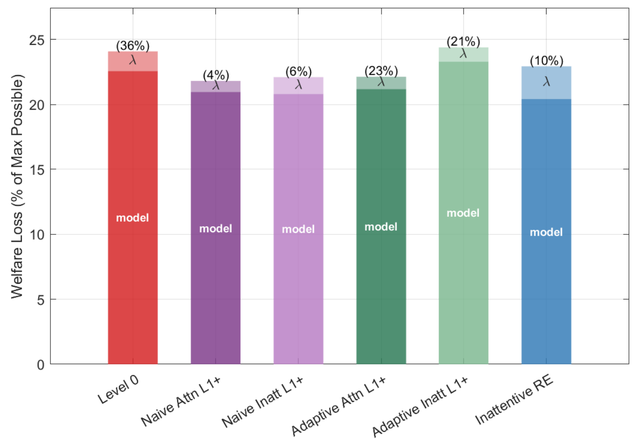

Figure 15: Mean welfare loss by best-fitting category

Notes: Welfare loss is computed as $\ell_n = 1 - S_n/(T - t_0)$, where S_n is the subject's total quadratic score from trial $t_0 + 1$ onward. Losses are decomposed into model and λ components. The darker portion of each bar reflects losses due to model errors, while the lighter portion reflects expected losses from implementation costs λ. Percentage at the top of each bar indicates the fraction of subjects in the category.

sponse is faster. The *Attentive* models have the characteristic hump-shaped response, as the level-0 belief incorporates post-shock ring realizations at rate γ. Under *Adaptive Inattentive*, both learning and inattention compound, but because these subjects are more strategically sophisticated, their response to a regime change is faster. In all cases, control costs dampen the IRFs, with higher λ values contributing more to the attenuation. Since the *Naive Attentive* subjects also have the relatively higher control costs, their IRF is the most dampened by these costs. Lastly, since the adaptive level-0 subjects dominate the sample, the population IRF is closes to their IRF and note that it is very close to the empirical IRF reported in the previous section.

6 Conclusion

We present results from a strategic estimation task in which participants have the option to use either model-based reasoning or empirical pattern recognition to form forecasts. We show that when actions are strategic, adjustment in response to large, visible changes in the fundamental state is noisy and very sluggish–more so than in experiments without strategic considerations. Strategic sophistication is dispersed and limited—more so than has been found in experiments from strategic games literature. Finally, experience-based forecasting seems much more empirically-relevant than model-based forecasting. Even in an environment where deductive reasoning is simple and people have all the information they need to compute

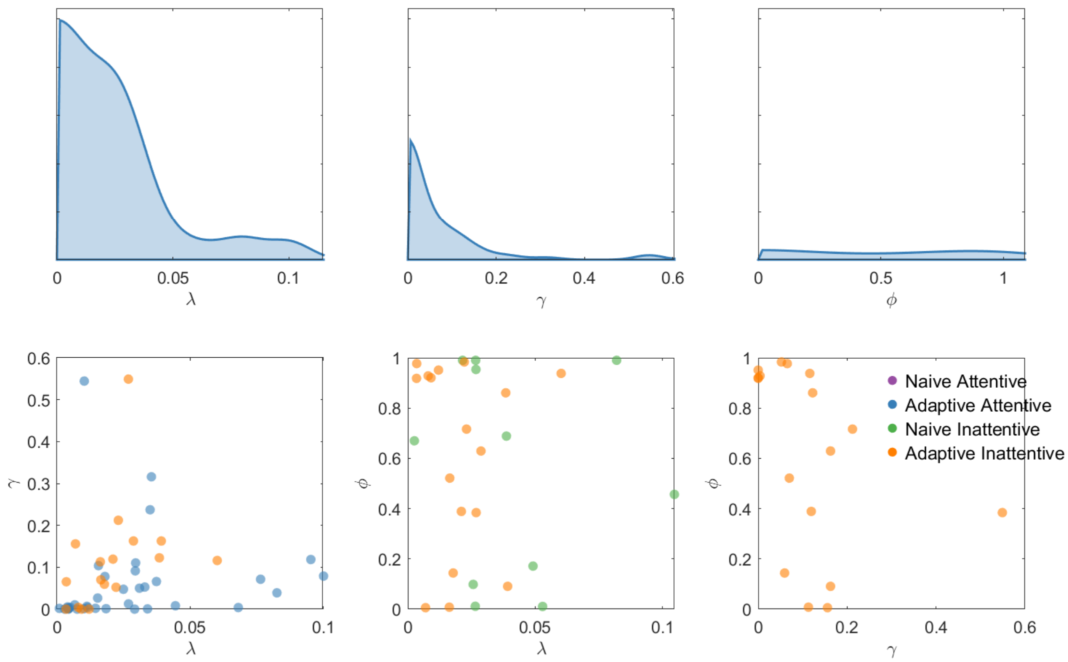

Figure 16: Parameter distributions and pairwise relationships. Top row: kernel density estimates of the marginal distributions of λ (effort cost), γ (learning gain), and ϕ (inattentiveness) across all subjects for whom each parameter is estimated. The y-axis scale is shared across panels. Bottom row: pairwise scatter plots colored by best-fitting model. Note: γ is estimated only for adaptive models and ϕ only for inattentive models, so the number of observations varies across panels.

the REE, not only do people not jump immediately to the new equilibrium, but there is a surprising amount of statistical learning. Since it yields very slow adjustment to regime changes, this finding has implications for policy changes. In practice, complications regarding the environment and potentially ambiguous announcements of policy changes are likely to make learning and extrapolation of patterns even more appealing than in the controlled environment of the lab.

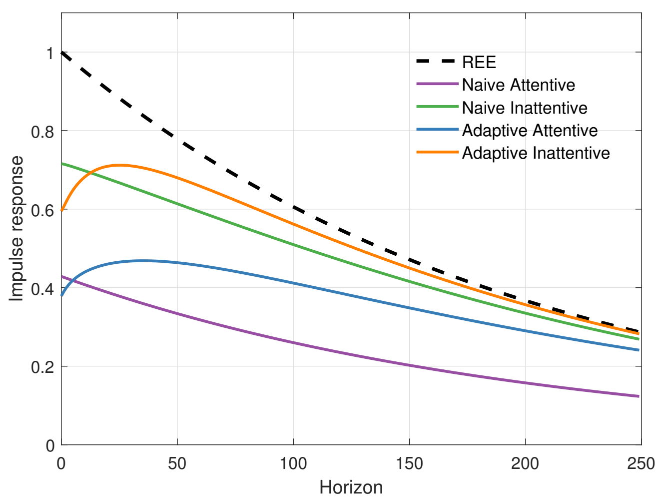

Figure 17: Model-implied impulse response functions

Notes: IRF for best-fitting models.

References

Calvo, Guillermo A. (1983), "Staggered prices in a utility-maximizing framework," *Journal of monetary Economics* 12(3): 383–398.

Carroll, Christopher D, Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka & Matthew N White (2020), "Sticky expectations and consumption dynamics," *American Economic Journal: Macroeconomics* 12(3): 40–76.

Coibion, Olivier & Yuriy Gorodnichenko (2015), "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review* 105(8): 2644–78.

Eusepi, Stefano & Bruce Preston (2018), "The science of monetary policy: An imperfect knowledge perspective," *Journal of Economic Literature* 56(1): 3–59.

Evans, George W & Seppo Honkapohja (1999), "Learning dynamics," *Handbook of Macroeconomics* 1: 449–542.

Evans, George W & Seppo Honkapohja (2013), "Learning as a rational foundation for macroeconomics and finance," *Rethinking expectations: The way forward for macroeconomics* 68.

Farhi, Emmanuel & Iván Werning (2019), "Monetary policy, bounded rationality, and incomplete markets," *American Economic Review* 109(11): 3887–3928.

Fehr, Ernst & Jean-Robert Tyran (2008), "Limited rationality and strategic interaction: the impact of the strategic environment on nominal inertia," *Econometrica* 76(2): 353–394.

Gallistel, Charles R., Monika Krishan, Ye Liu, Reilly Miller & Peter E. Latham (2014), "The Perception of Probability," *Psychological Review* 121(1): 96–123.

García-Schmidt, Mariana & Michael Woodford (2019), "Are low interest rates deflationary? A paradox of perfect-foresight analysis," *American Economic Review* 109(1): 86–120.

Kass, Robert E & Adrian E Raftery (1995), "Bayes factors," *Journal of the American Statistical Association* 90(430): 773–795.

Khaw, Mel Win, Luminita Stevens & Michael Woodford (2017), "Discrete adjustment to a changing environment: Experimental evidence," *Journal of Monetary Economics* 91: 88–103.

Mankiw, N Gregory & Ricardo Reis (2002), "Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve," *The Quarterly Journal of Economics* 117(4): 1295–1328.

Mauersberger, Felix & Rosemarie Nagel (2018), "Levels of reasoning in Keynesian Beauty Contests: A generative framework," in *Handbook of Computational Economics*, vol. 4, pp. 541–634, Elsevier.

McKelvey, Richard D & Thomas R Palfrey (1995), "Quantal response equilibria for normal form games," *Games and Economic Behavior* 10(1): 6–38.

McKelvey, Richard D & Thomas R Palfrey (1998), "Quantal response equilibria for extensive form games," *Experimental Economics* 1: 9–41.

Nagel, Rosemarie (1995), "Unraveling in guessing games: An experimental study," *The American Economic Review* 85(5): 1313–1326.

Ricci, M & C R Gallistel (2017), "Accurate step-hold tracking of smoothly varying periodic and aperiodic probability," *Attention, Perception, & Psychophysics* 79(5): 1480–1494.

Smith, Vernon L, Gerry L Suchanek & Arlington W Williams (1988), "Bubbles, crashes, and endogenous expectations in experimental spot asset markets," *Econometrica* pp. 1119–1151.

Stahl, Dale O & Paul W Wilson (1994), "Experimental evidence on players' models of other players," *Journal of Economic Behavior & Organization* 25(3): 309–327.

Stahl, Dale O & Paul W Wilson (1995), "On players' models of other players: Theory and experimental evidence," *Games & Economic Behavior* 10(1): 218–254.

Vimercati, Riccardo Bianchi, Martin S Eichenbaum & Joao Guerreiro (2021), "Fiscal policy at the zero lower bound without rational expectations," NBER Working Paper w29134.

Woodford, Michael (2003), "Imperfect Common Knowledge and The Effects of Monetary Policy," in *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, J. Stiglitz P. Aghion, R. Frydman & M. Woodford, eds., pp. 25–58, Princeton University Press, Princeton, NJ.