A Unified Model of Learning to Forecast*

George W. Evans
University of Oregon and University of St Andrews

Christopher G. Gibbs
University of Sydney

Bruce McGough
University of Oregon

August 30, 2023

Abstract

We propose a model of boundedly rational and heterogeneous expectations that unifies adaptive learning, k-level reasoning, and replicator dynamics. Level-0 forecasts evolve over time via adaptive learning. Agents revise over time their depth of reasoning in response to forecast errors, observed and counterfactual. The unified model makes sharp predictions for when and how fast markets converge in Learning-to-Forecast Experiments, including novel predictions for individual and market behavior in response to announced events. We present experimental results that support these predictions. We apply our unified approach in the New Keynesian model to study forward guidance policy.

JEL Classifications: E31; E32; E52; E71; D84; D83.

Key Words: adaptive learning and level-k reasoning; behavioral macroeconomics; sacrifice ratio; forward guidance; experiments.

*We thank Yuqi (Angela) Jiang and Laura Flakowski for research assistant support and Timothy Johnson for programming support. This paper has benefited from comments received from Valentyn Panchenko, Isabelle Salle, Yuta Takahashi and in talks at the University of Amsterdam, Universitat Rovira i Virgili, Australian National University, the University of Adelaide, University of St Andrews, Hitotsubashi University, the 26th Symposium of SNDE at Keio University, the 2018 Expectations and Dynamic Macroeconomic Models Conference hosted by the University of Birmingham, and the 2022 Barcelona Summer Forum Expectations and Dynamic Macroeconomic Models. We also wish to thank Colin Camerer who shared some experimental data with us at a very early stage. Financial support from National Science Foundation Grant No. SES-1559209, Australian Research Council under the grant DP210101204, UNSW Business School, and The University of Sydney is gratefully acknowledged.

1 Introduction

The assumption of rational expectations (RE) continues to come under scrutiny in macroeconomics and finance models in which RE plays a central role. RE imposes strong assumptions on agents' knowledge and cognitive abilities that call into question the plausibility and robustness of some model predictions. This issue is particularly acute when studying the general equilibrium implications of structural change in RE models in which there are several salient empirical puzzles, including the well known forward guidance puzzle.

Increasingly, modelers are turning to boundedly rational alternatives to RE such as adaptive learning (e.g. Evans, Honkapohja and Mitra, 2009 and Gibbs and Kulish, 2017), level-k reasoning (e.g. Angeletos and Lian, 2018, García-Schmidt and Woodford, 2019, and Farhi and Werning, 2019), and behavioral models (e.g. Arifovic, Schmitt-Grohé and Uribe, 2018 and Goy, Hommes and Mavromatis, 2020) to attempt to resolve the puzzles.[1] A common justification advanced by many of these studies is that there is ample evidence to support their modeling choices from laboratory experiments.

It is crucial to stress the equilibrium nature of RE, in most economic settings, which is seen most clearly in the simple guess-the-average game. Subjects are asked to pick a number between 0 and 100, in which the winning number is the one closest to 2/3 of the average guess. A subject who treats 50 as focal point, e.g. because it is the mean of a random guess from [0,100], might then choose the "level-1" guess of $(2/3) \times 50$. However a subject who thinks that other subjects make the level-1 guess, may make the level-2 guess of $(2/3)^2 \times 50$, etc. The unique Nash equilibrium guess is zero, but in one-shot games, even for grandmaster chess players, average guesses are typically over 30, with winning guesses over 20.[2] In many RE models, including Muth (1961), Lucas (1972), and the New Keynesian model, this type of strategic uncertainty is implicitly present, but less evident and often ignored, in the REE solution.

This paper seeks to unify key elements of alternative bounded-rationality approaches to strategic uncertainty by marrying adaptive learning and level-k reasoning within a single heterogeneous expectations behavioral model. Adaptive learning and heterogeneous expectations capture well-documented behavior of

[1] In the formal model that we analyze in detail, we assume, in line with most of the RE literature, that agents have full information about the structure of the economy. It is, of course, also possible to retain RE while relaxing the full information assumption, e.g. Bianchi, Lettau and Ludvigson (2022) and Bianchi and Melosi (2014) consider models in which agents are rational but face uncertainty about the duration of an observed regime change. In our macroeconomic policy example in Section 7, we illustrate how our approach can be extended to a setting in which there is uncertainty about the duration of an observed policy regime.

[2] See Nagel, Bühren and Frank (2017).

laboratory subjects in Learning-to-Forecast Experiments (LtFE), e.g. Hommes, Sonnemans, Tuinstra and Van De Velden (2007), Hommes (2011), and Hommes (2013).[3] At the same time, level-k reasoning provides a natural way to model the general equilibrium implications of forward-looking boundedly rational expectation formation, and it enjoys wide experimental support as shown by Nagel (1995), Duffy and Nagel (1997), Ho, Camerer and Weigelt (1998), Bosch-Domenech, Montalvo, Nagel and Satorra (2002), Costa-Gomes and Crawford (2006), Nagel et al. (2017) and Mauersberger and Nagel (2018). This literature has also shown that in repeated "guess the average" games, a special case of our benchmark univariate model, agents tend to shift over time towards higher level-k forecasts.

Our model is populated by agents with perfect knowledge of the structure of the economy, but imperfect knowledge of the expectations of others. To form forecasts, agents choose a sophistication level, k, that reflects level-k deductions along the lines of Nagel (1995). Specifically, there is a forecasting strategy of minimal sophistication, level-0, that uses a model-related salient value, which in our dynamic setting will be history dependent, adapting to observed data as discussed below. Level-1 agents use their knowledge of the economic environment to choose a forecast that would be optimal if all other agents are level-0; the forecasts of level-k agents are defined inductively.

Central to our approach is that agents (i) have knowledge about the economic structure, and specifically about how outcomes depend on the expectations of other agents, but (ii) cannot directly observe the expectations of other agents. Rational agents may thus understand that there is an RE equilibrium (REE), yet refrain themselves from holding expectations consistent with the REE because they doubt that other agents will hold RE.[4] The validity of these doubts is well-established in the above-cited experimental literature.

Of course, rational agents might well consider the possibility that other agents have heterogeneous expectations, in particular that other agents hold heterogeneous level-k beliefs. For any given distribution of level-k beliefs, the corresponding optimal expectation could be computed; however, it is implausible that an agent would know this distribution. An advantage of the level-k approach is that it focuses on choosing from an easily computable set of forecasts based on depth of reasoning: level-k expectations are optimal when the *average* expectation held by other agents is level k-1.

[3] LtFE are laboratory experiments in which the sole or principal task of the subject is to make forecasts of key economic variables.

[4] The eductive learning literature, e.g. Guesnerie (1992), Guesnerie (2002), emphasizes that both structural restrictions and strong higher-order common knowledge assumptions would be needed for fully rational agents to coordinate on the REE. Eductive reasoning more broadly refers to using knowledge of the economic structure to make inferences about the possible expectations held by other rational agents.

Also central to our approach is a dynamic setting in which agents, each period, make decisions based on their forecasts, and then observe outcomes. This allows agents to learn from the data over time in two distinct ways: First, an adaptive learning rule adjusts the level-0 forecast each period in response to observed outcomes. Second, agents engage in predictor selection, based on replicator dynamics. Level-k predictors that generate large forecast errors lose users to the best level-k forecasts.

In stationary environments adaptive learning is known to converge over time to rational forecasts, in a wide range of settings, and yet requires no knowledge of structural parameters. It thus provides a simple, robust, and natural way to model the evolution of level-0 forecasts. The motivation for the replicator dynamics goes to our observation that agents have no information on the distribution of different forecasts currently in use other than recent observations of actual outcomes; the most natural dynamic for the proportions of level-k forecasts is therefore for them to shift over time toward the k-level that would have provided the most accurate forecast.

To summarize, our bounded rationality model includes three elements:

1. Adaptive learning to modify level zero forecasts.
2. A menu of level-k forecasts computed using the known structure.
3. Replicator dynamics that shift agents towards the optimal reasoning level.

We establish important theoretical results, which include in particular that in stationary environments unified dynamics can generate rational expectations equilibria as emergent outcomes. We then use our approach to explain the findings of lab experiments, and to examine implications of structural change in our univariate model and policy change in our macroeconomic application.

We study the unified model first in the univariate set-up of Muth (1961). After deriving sharp analytical results and examining simulations that illustrate the model's implications for different types of expectational feedback, we take the model to the laboratory and test its core predictions using a standard experimental design. We then extend our framework to the New Keynesian model and reconsider the analysis of Bilbiie (2019) on optimal forward guidance. We show that the unified model justifies low the level-k assumptions adopted in prominent papers such as Angeletos and Lian (2018) and Farhi and Werning (2019).[5]

There are other related approaches in the literature. One strand is to assume fixed proportions of agents that differ in their degree of sophistication. In the simplest cases there are just two categories, unsophisticated agents that use a non-rational forecast rule and other agents that are fully rational, taking into

[5]In fact, our model explicitly addresses limitations in the existing literature noted by Farhi and Werning (2019): see their quote we provide in Section 7.5 below.

account the proportion of nonrational agents. For example Gali and Gertler (1999) and Jackson (2005) consider inflation dynamics when a fixed proportion of agents follow naive rule-of-thumb forecasts, while the other agents have rational expectations that take into account the naive forecasters. Mokhtarzadeh and Petersen (2021) explore a monetary model in which a proportion of agents have expectations aligned with central bank forecasts, while the other type have RE in the sense that they best-respond to the first type.

However, the models just described assume additional knowledge of the sophisticated agents beyond knowing the structure of the economy: the sophisticated agents must know the proportion of unsophisticated agents and the specific forecast rules those agents are following. Of even greater concern to us is that these approaches do not address the strategic uncertainty that underlies the beliefs of the rational agents: optimal decisions by agents depend on the prices they expect, but those prices depend on the expectations of other agents, which, if these other agents include rational agents, depend on the expectations that other rational agents expect other rational agents to hold, ad infinitum. Truly sophisticated agents, who know the economic structure, will not align their expectations with RE if they are concerned that other sophisticated agents, at some level of this recursion, do not have "rational" expectations.

Our framework specifically addresses this concern and shows that the extent to which agents coordinate on rational expectations can evolve over time through both adaptive and "eductive" level-k channels. It would be possible to extend our model to include a known fixed proportion of naive agents that follow a known specified rule-of-thumb, or to include a proportion of agents that are fully rational in that they coordinate on rational expectations given the proportions and forecasts of all level-k (and naive) forecasts.[6]

The Cognitive Hierarchy (CH) approach of Camerer, Ho and Chong (2004), in contrast, allows for a distribution of "k-step" types. In the CH framework this distribution satisfies two assumptions. First, every agent believes, incorrectly, that there are no other agents with equal or higher k-step beliefs; second, every agent knows the exact relative distribution of lower k-step agents. Given these beliefs, k-step agents make optimal decisions conditional on the implied forecasts obtained from those beliefs. Camerer et al. (2004) focus on a family of Poisson distributions that satisfy these assumptions. From our perspective, it is difficult to understand how agents could come to know the distribution of lower k-step

[6] In an earlier version of the paper we included a proportion of agents that were fully rational, possessing the knowledge just stated, and we provided the "eductive stability" condition needed for the rational agents to achieve expectational coordination under suitable, strong common knowledge conditions along the lines discussed in Guesnerie (2002). This is omitted in the current paper because of the additional complexity entailed.

types, yet at the same time not realize or consider that there are other agents using equal or higher reasoning steps. These concerns would appear to be even more acute in extensions to repeated or dynamic games.

In the CH approach, as well as in the Reflective Equilibrium approach of García-Schmidt and Woodford (2019) discussed in Section 8, the analysis takes place at a single point in time. In contrast our unified model provides full real-time dynamics for the time paths of level-k forecasts, the proportions of agents using each forecast level, and for the associated time path of the endogenous variables. Our model is particularly suited for analysis of the impact over time of announced future policy changes, an issue that has received considerable attention in the macroeconomic policy literature.

Because our model integrates distinct boundedly rational approaches to decision making over time, we are drawing on many earlier works. Section 8 locates our paper within this broader literature.

2 Overview of Model and Results

We develop our unified approach using the benchmark univariate linear "cobweb" model of Muth (1961), allowing for either positive or negative feedback. After fully examining the univariate set-up, both theoretically and experimentally, we show how to extend this framework to a multivariate forward-looking New Keynesian model and examine its implications for announced policy changes including, in particular, monetary policy forward guidance.

The univariate model takes the form $y_t = \gamma + \beta \hat{E}_{t-1} y_t$, where $\hat{E}_{t-1} y_t$ is the average of individual forecasts, made at time $t-1$, of the variable y_t. Assume $\beta \neq 0, 1$, and for simplicity assume there are no exogenous random shocks so that the REE is $y_t = \bar{y} = (1-\beta)^{-1} \gamma$. The case $\beta < 0$, with negative expectational feedback, corresponds to the Muth cobweb model of prices in an isolated market with a production lag, while the case $0 < \beta < 1$, with positive expectational feedback, corresponds to a repeated beauty contest or guess-the-average game.

Agents have heterogeneous level-k forecasts. Letting $a_{t-1} = E^0_{t-1} y_t$ denote the level-0 forecast at $t-1$, level-k forecasts are defined recursively:

$$E^k_{t-1} y_t = \gamma + \beta \left(E^{k-1}_{t-1} y_t \right), \text{ for } k = 1, 2, 3 \ldots$$

Letting ω_{kt} denote the proportion of agents with level-k forecast $E^k_{t-1} y_t$, we have $y_t = \gamma + \beta \sum_{k \geq 0} \omega_{kt} E^k_{t-1} y_t$. Under our unified approach, there are two channels of learning dynamics. First, level-0 dynamics are driven by standard adaptive learning rules updating a_t toward the most recent observation y_t. Second, weights ω_{kt} are updated each period based on replicator dynamic that shifts weight toward the k level providing the most accurate forecast the previous period.[7]

Section 3 gives the formal details of the model, including the adaptive and replicator mechanisms that generate the unified dynamics. Section 4 presents a formal analysis of the asymptotic properties of the unified model, together with quantitative illustrations of how qualitative features of the dynamics depend on the feedback parameter β, both in a stationary setting and in response to announced structural changes.

When $\beta > 1$ the asymptotic dynamics in a stationary setting are unstable, so that $|y_t| \to \infty$, whereas if $|\beta| < 1$ there is convergence over time to the REE under unified dynamics, i.e. $y_t \to \bar{y}$. When $|\beta| < 1$ the adaptive learning and the replicator mechanisms are each sufficient to deliver asymptotic convergence. If adaptive learning is shut down the replicator dynamic generates asymptotic convergence by shifting weights over time to higher k-levels. If instead the distribution of k-levels is fixed over time, $\omega_{kt} = \omega_k$, the adaptive learning dynamics induces convergence of level-0 forecasts to the REE. Thus when $|\beta| < 1$ adaptive and level-k replicator dynamics are complementary. In contrast, when $\beta < -1$ adaptive learning and level-k dynamics can work against each other.

Simulations of the unified model provide additional insights. When $|\beta| < 1$ convergence can be much faster than under adaptive or replicator dynamics alone, and can lead to a mixture of high and low-level reasoners for extended periods with $y_t \approx \bar{y}$. When $\beta < -1$ our model makes other novel predictions: convergence to the REE, unstable dynamics, and bounded cycles that are not centered at the REE are all possible. These findings can provide an explanation for the experimental results of Bao and Duffy (2016) that when $\beta < -1$ market dynamics are distinctly different: they observe both stable and unstable cases.

Sections 5 and 6 take our model to the lab.[8] In our LtFE, we adapt the experimental design of Bao and Duffy (2016) to test key predictions of the unified model. We place laboratory subjects into a computer-based market that nests the cobweb model. Participants have full information of the market structure. We consider both positive and negative expectational feedback cases. A novel dimension of our experiment is announced structural changes at irregular intervals, which allows us (and the participants) to clearly identify the level-0 beliefs. The unified model then provides sharp predictions for the distribution of forecasts

[7] Our framework could be generalized in various ways. The level-k menu is straightforward to compute, and thus serves to provide natural focal points for forecasts. However, the menu could be extended to include, for example, the average of level-k and level-$k+1$ forecasts. One could also include level-k calculation costs that are increasing in k. Although extensions like this may prove fruitful in experimental or empirical work, we conjecture that our central findings would be qualitatively unaffected.

[8] Readers wanting to focus on the macroeconomic application can move directly from Section 4 to Section 7 without loss of continuity.

observed in announcement periods, as well as for how people should revise their depth of reasoning in subsequent periods.

We find strong evidence for both adaptive and level-k type reasoning underlying expectations. In particular, in announcement periods we can classify between 50% to 70% of participants, depending on the measure, as either level-0, 1, 2, 3 or as those who use a value close to the REE forecast. Moreover, we find that larger numbers of subjects are classified as playing k-level strategies in later announcement rounds. In our experiment level-k behavior is observed across all treatments and is particularly prominent when $\beta < 0$. In this latter case, we observe subjects making clear level-k deductions that oscillate above and below the perfect foresight equilibrium, behavior that is sometimes argued to be implausible when level-k reasoning is adapted to more complex macroeconomic environments as in García-Schmidt and Woodford (2019) and Angeletos and Sastry (2021).

We also find evidence for an additional prediction of the unified model: in the wake of announcements subjects may lower their depth of reasoning. We document, as the model predicts, that some high-level reasoners experience large forecast errors in announcement treatments. This causes a fraction of the high-level reasoners to revise down their depth of reasoning. These downward revisions can make the prevalence of low-level reasoning very persistent

Section 7 extends the unified framework to a multivariate forward-looking macroeconomic model with an announced policy change. Specifically, for New Keynesian model

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^n), \;\; \pi_t = \xi E_t \pi_{t+1} + \kappa x_t,$$

modified to include unified dynamics, we consider monetary policy forward guidance following a persistent stagnation shock, as in Bilbiie (2019).[9] We find that the coupling of adaptive learning and replicator dynamics endogenously induces low level reasoning, substantially reducing the power of monetary policy promises.

Section 8 discusses related literatures providing foundations for our unified model as well as related theoretical and experimental results. Section 9 concludes.

3 The Model

In this section we develop the static version of the model, which includes agents with varying levels of forecast sophistication. Incorporating dynamics via two distinct mechanisms through which agents can improve their forecasts over time, we present and analyze the unified model.

[9] Here, as usual, x is the output gap, π is inflation, and i and r^n are the nominal and natural, real interest rates, respectively.

3.1 The static model

There is a continuum of agents. The aggregate variable at time t, given by y_t, is determined entirely by the expectations of these agents, who are partitioned into a finite number of types. Types are distinguished by sophistication level, which is naturally indexed by the non-negative integers $\mathbb{N}$. For $k \in \mathbb{N}$, the proportion ω_k of agents of type k (i.e. having sophistication k) is referred to as the *weight* associated with agent-type k. The distribution of agents across types is summarized by a *weight system* $\omega = \{\omega_0, \ldots, \omega_M\}$, which is a vector of non-negative real numbers that sums to one, and where M is the number of agent types, which, in our dynamic settings, will typically be endogenously determined and vary over time.[10] We denote by Ω the collection of all possible weight systems as M varies over $\mathbb{N}$. This set, together with its natural topology, will be relevant for some of the analytic work in Section 4.

The forecasts made by agents with sophistication level k is given by $E^k_{t-1}y_t$, where higher k indicates greater sophistication. Aggregate y_t is determined as

$$y_t \;=\; \gamma + \beta\textstyle\sum_{k=0}^{M}\omega_k E^k_{t-1}y_t \equiv \gamma + \beta\sum_k \omega_k E^k_{t-1}y_t, \qquad (1)$$

where the equivalence on the right emphasizes that the implicitly limited sum ranges over the indices of the given weight system, a convention we adopt throughout the paper. We assume that $\beta \neq 0, 1$, and note that equation (1) nests the beauty contest or guess-the-average game, as well as the cobweb model. We note also that there is a unique equilibrium $\bar{y} = \gamma(1-\beta)^{-1}$ in which all agents have perfect foresight: this equilibrium corresponds to the rational expectations equilibrium (REE) of the simple RE model $y_t = \gamma + \beta E_{t-1}y_t$.

In our set-up, greater sophistication solely reflects higher order beliefs, as in the level-k framework of Nagel (1995). Agents with *level*-0 *beliefs* hold a common prior and form their forecasts accordingly as $E^0_{t-1}y_t = a$. Agents with higher-order beliefs are assumed to have full knowledge of the model. We recursively define *level-k beliefs* as beliefs that would be optimal if all other agents used level $k-1$:

$$E^1_{t-1}y_t = T(a) \equiv \gamma + \beta a \text{ and } E^k_{t-1}y_t = T^k(a) \equiv T\left(T^{k-1}(a)\right) \text{ for } k \geq 2.$$

Note that for $k \geq 1$ agents are assumed to know β and γ.[11]

The most natural level-0 belief will depend on the model. For example, the level-0 belief may reflect a salient value, as in the guessing game model in Nagel

[10] Sometimes we view a weight system as a sequence with only finitely many nonzero terms.

[11] This assumption makes modeling anticipated changes, like those implemented in our experiments, straightforward: any changes to β or γ known at time $t-1$ that occurs in time t are built directly into the forecasts of agents for which $k \geq 1$.

(1995) where this is taken as the midpoint of the range of possible guesses; or, in the cobweb model, the level-0 belief might be determined by the previous equilibrium in a market-setting, before a structural change has occurred, or it may be determined adaptively by looking at past data.

Combining these definitions with equation (1) yields the realized value of y as a function of level-0 beliefs, i.e. $y_t = \mathcal{T}(a)$, where

$$\mathcal{T}(a) = \gamma \left(1 + \frac{\beta}{1-\beta} \sum_{k \geq 0} \left(1 - \beta^k\right) \omega_k \right) + \left(\beta \sum_{k \geq 0} \beta^k \omega_k \right) a. \tag{2}$$

We note that $\mathcal{T}$ is linear in a, and it is convenient to rule out the non-generic case that the coefficient on a, given by, $\beta \sum_{k \geq 0} \beta^k \omega_k$, has a modulus of one. Finally, we remark that the REE is a fixed point of $\mathcal{T}$, i.e. $\mathcal{T}(\bar{y}) = \bar{y}$.

It would be possible to extend the model to include a class of agents who are fully rational, which, in our environment, would correspond to perfect foresight. This would require rational agents to fully understand the distribution and behavior of all agent types. In the current setting this appears implausible and, at the same time, would lead to further complexity. For example, the inclusion of rational agents requires additional stability considerations to ensure coordination of the rational agents, given the forecasts of the other agents. The appropriate condition is the eductive stability condition when the economy includes non-rational agents and is given in Gibbs (2016).

3.2 Adaptive dynamics

We define *adaptive dynamics* as corresponding to adaptive learning with fixed level-k weights.[12] Specifically, a weight system ω is taken as fixed and level-0 forecasts $E^0_{t-1} y_t \equiv a_{t-1}$ are assumed to evolve over time in response to observed outcomes. The system under adaptive dynamics is given by

$$\begin{aligned} y_t &= \gamma + \beta \textstyle\sum_{k \geq 0} \omega_k E^k_{t-1} y_t \\ a_t &= a_{t-1} + \phi(y_t - a_{t-1}), \end{aligned} \tag{3}$$

where $0 < \phi < 1$. The simple form of the updating rule for level-0 beliefs reflects that our model is univariate and non-stochastic. The parameter ϕ, often called the gain parameter, specifies how much the forecast adjusts in response to the most recent forecast error. The time t forecasts a_t can be equivalently written as a geometric average of previous observations with weights $(1-\phi)^i$ on y_{t-i},

[12] We use the term "adaptive dynamics" to distinguish our model and results from the well-understood "adaptive learning" case in which all agents are level-0.

for $i \geq 1$.[13] Backward-looking rules like (3), as well as anchor and adjustment rules and trend following rules, are frequently found to well-describe behavior of laboratory participants in LtFEs as discussed in Hommes (2013). We focus on the specification (3) in order to emphasize the novel features of our framework.

3.3 Replicator dynamics

We next consider the possibility that agents revise their depth of reasoning over time based on their past forecast performance. Nagel (1995) and Duffy and Nagel (1997) each explore whether laboratory participants update their depth of reasoning over time in repeated guess-the-average experiments. They find that in general they do not update their reasoning in games with few repetitions - four or fewer - but do appear to do some updating in games of 10 repetitions or more. To capture this sort of updating behavior, we consider the possibility that agents are relatively inattentive to revising their depth of reasoning. More specifically, we assume that (typically) only a small proportion of agents using sub-optimal reasoning levels will revisit and revise their forecast methods, with the proportion begin dependent on forecast error magnitude. This captures the behavioral premise of Kahneman (2011) that much of decision-making is based on "thinking fast" routinized procedures (in our case, using the same forecast method as in the previous period), while larger errors incline more agents to "think slow," (in our case, revisit and revise their reasoning depth).

We formalize this process by appealing to a kind of replicator dynamic along the lines of those considered by Weibull (1997), Sethi and Franke (1995), and Branch and McGough (2008). We assume the best level-k forecast gains more users over time while more poorly performing forecasts lose users over time. Importantly, the largest depth of reasoning considered is endogenous: agents are allowed to consider reasoning depths that have never been played in the game.

The replicator dynamic we propose shifts weight from suboptimal predictors towards the (time-varying) optimal predictor according to a "rate" function that depends on the forecast error. We define the time t optimal predictor as

$$\hat{k}(y_t) = \min \arg\min_{k \in \mathbb{N}} |E^k_{t-1} y_t - y_t|, \tag{4}$$

where the left-most "min" is used to break ties.[14]

[13] If y_t, in equation (1), also depended on a white-noise random shock then ϕ would typically be replaced by a time-varying term that decreases asymptotically at rate $1/t$. In cases where y_t also depends on observable exogenous stochastic shocks, adaptive learning is formulated in terms of recursive least-squares updating. We conjecture – and provide experimental evidence in Section 6 – that some heterogeneity in the level-0 agents' learning rules, or some perceived heterogeneity of the rule by $k \geq 0$ types, would not materially affect our conclusions.

[14] Existence of $\hat{k}(y_t)$ is assured: $k \to \infty$, $|\beta| < 1$ ($|\beta| > 1$) implies $E^k_{t-1} y_t \to 0$ ($E^k_{t-1} y_t \to \infty$).

Next, assume the presence of a rate function $r : [0, \infty) \to [\delta, 1)$ with $\delta \geq 0$ satisfying $r' > 0$.[15] Finally, let ω_{kt} be the weight of level-k beliefs in period t. The system under replicator dynamics is given by period t according to

$$\begin{aligned} y_t &= \gamma + \beta \sum_{k \geq 0} \omega_{kt} E^k_{t-1} y_t \\ \omega_{it+1} &= \begin{cases} \omega_{it} + \sum_{j \neq \hat{k}(y_t)} r\left(\left|E^j_{t-1} y_t - y_t\right|\right) \omega_{jt} & \text{if } i = \hat{k}(y_t) \\ \left(1 - r\left(\left|E^i_{t-1} y_t - y_t\right|\right)\right) \omega_{it} & \text{else} \end{cases} \end{aligned} \tag{5}$$

We note that the replicator dynamic requires a given value a for level-0 beliefs, as well as an initial weight system $\omega_0 = \{\omega_{k0}\}_{k \in \mathbb{N}}$.

3.4 Unified dynamics

Unified dynamics joins adaptive dynamics and replicator dynamics. The level-0 forecasts are updated over time as in Section 3.2 and the weights evolve according to the replicator as in Section 3.3. The system under unified dynamics is given as

$$\begin{aligned} y_t &= \gamma + \beta \sum_{k \geq 0} \omega_{kt} E^k_{t-1} y_t = \gamma + \beta \sum_{k \geq 0} \omega_{kt} T^k\left(a_{t-1}\right) \\ \omega_{it+1} &= \begin{cases} \omega_{it} + \delta_r \sum_{j \neq \hat{k}(y_t)} r\left(\left|T^j\left(a_{t-1}\right) - y_t\right|\right) \omega_{jt} & \text{if } i = \hat{k}(y_t) \\ \left(1 - \delta_r r\left(\left|T^k\left(a_{t-1}\right) - y_t\right|\right)\right) \omega_{it} & \text{else} \end{cases} \\ a_t &= a_{t-1} + \phi(y_t - a_{t-1}), \end{aligned} \tag{6}$$

where $\delta_r \in \{0, 1\}$ indicates whether the replicator dynamic is operable. We note that while the adaptive dynamics and replicator dynamics can be viewed as special cases of the unified model, it is useful (and even necessary) to analyze them in isolation; and we proceed this way in the next section.

Our interests include the economy's asymptotic properties. We say the *model* is *stable* if y_t converges to the perfect foresight equilibrium $\bar{y}$ for all relevant initial conditions, which, in case of the unified dynamic, include initial beliefs a and initial weights ω. We say the *model* is *unstable* if $|y_t| \to \infty$ for all relevant initial conditions, with $a \neq 0$. We will find that stability and instability can be fully characterized when $\beta > -1$, but that with large negative feedback there is a tension between stabilizing and destabilizing forces.

4 Properties of the unified model

In this section we develop the analytic properties of the unified model. We begin by establishing the available analytic results, and the turn to simulations for

[15] An example of a suitable rate function is $r(x) = {}^2\!/\!_\pi \tan^{-1}(\alpha x)$, with $\alpha > 0$ providing a tuning parameter. We use this rate function for our simulation exercises.

additional insights. These insights are aided by some partial analytic results on the dependence of $\hat{k}$ on the feedback parameter β. In the dynamic setting, $\hat{k}$ determines how the depth of reasoning of agents changes over time.

4.1 Stability results

Adaptive dynamics and replicator dynamics are special cases of this model (with $\delta_r = 0$ or $\phi = 0$, respectively), in which additional insights are available; however, our central result concerns the stability of the unified model.[16]

Theorem 1 (Stability of unified dynamics)**.** *Assume $\delta_r = 1$ and $0 < \phi \leq 1$.*

1. *If $|\beta| < 1$ then the model is stable: $y_t \to \bar{y}$.*
2. *If $\beta > 1$ then the model is unstable: $|y_t| \to \infty$.*

We remark that if $\beta < -1$ then odd levels of reasoning introduce negative feedback while even levels result in positive feedback. These countervailing tendencies can result in interesting and complex outcomes; but they also make $\beta < -1$ difficult to analyze. For partial results under adaptive dynamics: see Theorem 3.

We turn now to the replicator dynamic with the adaptive learning mechanism shut down, i.e. $\phi = 0$. In this case we start from an arbitrary (non-zero) level-0 forecast that remains unchanged, and convergence takes place through the replicator dynamic shifting weights over time to more sophisticated, i.e. higher level, forecasts. We have the following result.

Theorem 2 (Stability of replicator dynamics)**.** *Assume $\delta_r = 1$ and $\phi = 0$.*

1. *If $|\beta| < 1$ then the model is stable: $y_t \to \bar{y}$. Also, $t \to \infty$ implies $\hat{k} \to \infty$ and $\omega_{kt} \to 0$ for all $k \geq 0$.*
2. *If $\beta > 1$ then the model is unstable: $|y_t| \to \infty$.*

Intuitively, when $|\beta| < 1$ the map $\mathcal{T}(a)$ operates as a contraction, and as a result the optimal forecast level is higher than the average level used by agents. This tends to shift weight under the replicator to increasingly higher levels over time. However, as will be seen in the simulations, the dynamics of ω_{kt} for any given level k can be non-monotonic and complex.

When the replicator is shut down, i.e. $\delta_r = 0$, some additional notation is needed. Denote the *n-simplex* by $\Delta^n \subset \mathbb{R}^{n+1}$,

$$\Delta^n = \left\{ x \in \mathbb{R}^{n+1} : x_i \geq 0 \text{ and } \sum_i x_i = 1 \right\}.$$

The earlier-defined set of all weight systems, Ω, is the disjoint union of these simplexes: $\Omega = \dot{\bigcup}_n \Delta^n$, where the dot over the union symbol emphasizes that,

[16] Proofs of all theorems and propositions are found in the online appendix A1.

as subsets of Ω, the Δ^ns are pairwise disjoint. The set Ω inherits a natural topology, sometimes called the *final topology*, from the relative topologies on the Δ^ns: $W \subset \Omega$ is open if and only if $W = \dot{\cup}_n W_n$, with $W_n \subset \Delta^n$ open in Δ^n.[17]

Using this notation, and given $\beta \in \mathbb{R}$, we may define $\psi_\beta : \Omega \to \mathbb{R}$ by $\psi_\beta(\omega) = \beta \sum_k \beta^k \omega_k$, which, we recall from (2), is the coefficient of a in the formulation of the map $\mathcal{T}$. The following theorem establishes results under adaptive dynamics.

Theorem 3 (Stability of adaptive dynamics)**.** *Suppose $\delta_r = 0$ and $0 < \phi \leq 1$.*

1. *If $|\beta| < 1$ then the model is stable: $y_t \to \bar{y}$.*
2. *If $\beta > 1$ then the model is unstable: $|y_t| \to \infty$.*
3. *If $\beta < -1$ then ψ_β is surjective, and*
   - *(a) If $\psi_\beta(\omega) > 1$ then the model is unstable: $|y_t| \to \infty$.*
   - *(b) If $1 - 2\phi^{-1} < \psi_\beta(\omega) < 1$ then the model is stable: $y_t \to \bar{y}$.*
   - *(c) If $\psi_\beta(\omega) < 1 - 2\phi^{-1}$ then model is unstable: $|y_t| \to \infty$.*
   - *(d) There exists open subsets Ω_s and Ω_u of Ω such that i) if $\omega \in \Omega_s$ then the model is stable: $y_t \to \bar{y}$. ii) If $\omega \in \Omega_u$ then the model is unstable: $|y_t| \to \infty$. iii) The complement of $\Omega_s \cup \Omega_u$ in Ω is nowhere dense, i.e. its closure has empty interior.*

Some comments are warranted. Items one and two of this theorem are analogous to the results obtained in Theorems 1 and 2; however, here we also ca draw conclusions when $\beta < -1$. The surjectivity of ψ results from the expanding magnitudes and oscillating signs of the β^n. Adaptive dynamics may be written

$$a_t = \text{constant} + (1 - \phi(1 - \psi))a_{t-1},$$

so that the surjectivity of ψ implies that stability and instability may obtain for any value of ϕ. From item 3(b), two additional conclusions can be immediately drawn, and we summarize them as a corollary:

Corollary 1. *Suppose $\delta_r = 0$ and $\beta < -1$.*

1. *If $-1 < \psi_\beta(\omega) < 1$ then the model is stable for all $0 < \phi < 1$.*
2. *If $\psi_\beta(\omega) < -1$ then the model is stable for sufficiently small $\phi > 0$.*

Finally, item 3(d) evidences the challenge of predicting outcomes under unified dynamics or replicator dynamics when $\beta < -1$. The stable and unstable collections of weight systems are open and effectively cover Ω; as weight systems evolve over time it is very difficult to determine whether they eventually remain in either the stable or unstable regions.

[17] The final topology on a disjoint union of topological spaces is the *direct limit* topology induced by the inclusion maps $\Delta^n \hookrightarrow \Omega$.

4.2 Some results on $\hat{k}$

The behavior of the replicator dynamic is determined by the optimal level of reasoning, $\hat{k}$. To gain intuition for the mechanics of the replicator, in this section we study the dependence of $\hat{k}$ on β for the special case of uniform weights. In the online Appendix we show that $\hat{k} = \hat{k}(\beta, \omega)$ is independent of a and γ.

Proposition 1 (Optimal forecast levels)**.** *Let $K \geq 1$ and $\omega^K = \{\omega_n\}_{n=0}^{K}$ be a weight system with weights given as $\omega_n = (K+1)^{-1}$. Let $\hat{k} = \hat{k}\left(\beta, \omega^K\right)$.*

1. *If $|\beta| < 1$ then $K \to \infty \implies \hat{k} \to \infty$ and $\hat{k}/K \to 0$.*

2. *For given K, (a) $\beta \to -1^- \implies \hat{k} \to \begin{cases} 1 \text{ if } K \text{ is even} \\ 0 \text{ if } K \text{ is odd} \end{cases}$*

   (b) $\beta \to -1^+ \implies \hat{k} \to \infty$.

Although Proposition 1 examines only the specific case of uniform weights, it reveals how contrasting results for the optimal choice of k depend on β. When $|\beta| < 1$ and K is large, an approximately optimal forecast can be achieved with k-level increasing in, but small relative to, K. However, with $\beta < -1$, but $|\beta|$ not too large, the optimal k takes values in $\{0, 1\}$, with the specific value determined by *aggregate parity*, which is an aggregate measure of optimism and pessimism.[18]

4.3 Simulated Dynamic of the Unified Model with Announcements

A novel feature of the unified model is that boundedly rational agents can respond to anticipated events by incorporating information about changes in the economic environment. To illustrate this feature of the unified model, we simulate an economy with a non-zero REE, $\bar{y} > 0$ and, assuming free disposal, a non-negativity constraint on y and $E_{t-1}^{k} y_t$, which will mirror our lab experiment discussed in the next section.

We assume that γ, the intercept in equation (1), undergoes two announced changes, which shifts the steady state REE of the economy. Each simulation is 50 periods with $\gamma = 60$ for $t < 20$, $\gamma = 90$ for $20 \leq t < 45$, and $\gamma = 45$ for $t \geq 45$. The agents know the structure of the economy, the announced changes, and take into account that $y_t \geq 0$ when making their forecasts following level-k depths of reasoning.[19] The announcements are spaced such that the economy has converged to the pre-change steady state $\bar{y}$, which then constitutes the level-0 forecast when the announced change takes place.

[18] Proposition 1′, in the online Appendix, provides further results.

[19] The timing of expectations in the model is $E_{t-1}y_t$: knowledge of the change is only relevant for the forecast in the period before it occurs. Section 7 examines the unified dynamic to models in which y_t depends on $E_t y_{t+1}$.

Figure 1 shows the simulated results for the unified dynamics for three different β's corresponding to the regions of interest identified by our stability theorems. The parameter choices, announcements, and simulation length exactly mirror our experimental setup detailed in the next section. Each row of figures corresponds to a different feedback setting. The first plot in each row shows the proportion of agents using the level-0, 1, 2, and 3 predictors. The second plot in each row shows the optimal predictor in use in each period. The third plot in each row shows the level-0, 1, 2, and 3 predictions in each period. The fourth plot in each row shows the equilibrium dynamics of y_t.

Starting with the $\beta = 0.5$ simulation, we note three features of the unified dynamics. First, despite the fact that $y_t = \bar{y}$ for many periods prior to the announcements, the model does not predict instantaneous convergence to the new REE in these periods. In other words, convergence to the REE does not imply REE predictions going forward. This is because when the market has converged, low-level reasoning forecasts provide similar predictions to the REE forecast, so a mass of low-level reasoners remains even after the market has converged. The existence of these low-level reasoners implies that the optimal depth of reasoning in the announcement period is also relatively low, in line with Proposition 1 (second panel, bottom row). This leads to large forecast errors for those using higher depths of reasoning. Second, in response to these large forecast errors, some high-level reasoners will revise their beliefs *down* to lower levels of reasoning (see the first and second panels, bottom row). This kicks off another transition period, where it takes time for the market to re-converge. And third, although agents revise down their depth of reasoning, the proportion who are using a high depth of reasoning remains greater than in the initial periods because not all agents revise their forecasting strategy each period (see first panel, bottom row, and recall that the proportion using $k > 3$ is not shown).

The top row of Figures 1 shows the simulation for $\beta = -0.9$. A sizable proportion of agents uses relatively low levels of deduction even though the economy has converged prior to the announcement. Therefore, in the announcement period, the optimal depth of reasoning is low. The announcements cause those using higher levels of deduction to make large forecast errors. Some proportion of the high-level reasoners then revise their depth of reasoning lower as a result.

Similar dynamics are found for a wide range of parameters with $|\beta| < 1$. The presence of low-level reasoners when the announcements occur triggers the dynamics shown in Figure 1. However, the mass of high-level reasoners generally increases over time with repeated announcements.

Figure 1: Unified dynamics with announced structural change in period 20 and 45.

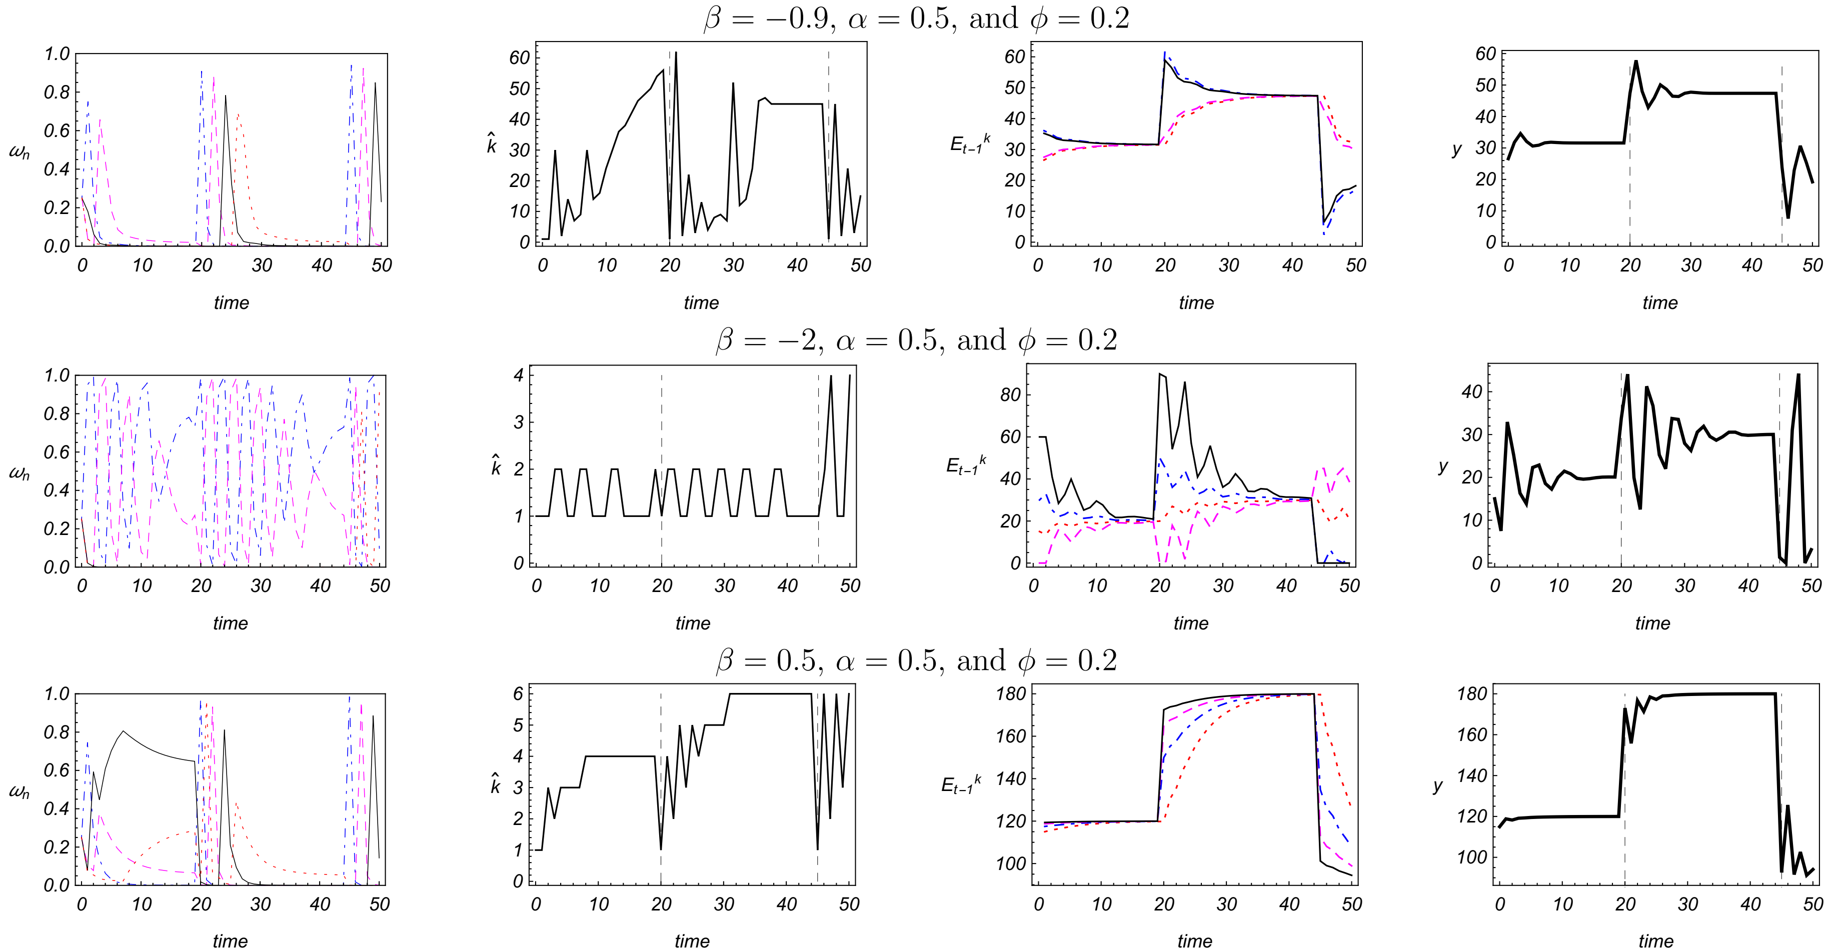

Notes: Simulation of unified dynamics with announced changes to the intercept and a known non-negativity constraint. $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *red dotted, blue dash-dot, dashed magenta* and *solid black*, respectively. The corresponding forecasts, $E^k_{t-1}y_t$, use the same style format.

The middle row of plots in Figure 1 shows a simulation for $\beta = -2$. Here the choice of parameters matters greatly for the outcome, and we consider a case in which the market converges after the first announcement. In contrast to the $|\beta| < 1$ cases, the optimal depths of reasoning do not rise over time. In fact, in order to stabilize the market, agents must choose relatively low depths of reasoning when y_t is not close to steady state. When y_t is away from the steady state, high depths of reasoning cause the non-negativity constraint to bind and predictions are either zero or γ. Therefore, the average depth of reasoning must remain low, in contrast to the previous cases, or y_t does not converge.

5 Learning-to-Forecast Experiment Design

The unified model makes distinctive predictions for individual expectations and market dynamics. To test these predictions we conduct a standard LtFE experiment following Bao and Duffy (2016). The experiment mirrors the simulated environment of Section 4.3 by having subjects participate in a repeated market for 50 periods, or *rounds.* They are asked to forecast the price of a good and they are compensated for the accuracy of their predictions. Market price is given by

$$p_t = \gamma + \beta \hat{E}_{t-1} p_t + \epsilon_t,$$

where $\hat{E}_{t-1} p_t$ is the average price forecast across participants and ϵ_t is a small white noise shock that is added to the system, which is standard practice in LtFE experiments. The shock sequence is the same in all markets and treatments.

We adopt a 3×3 experimental design where the treatment variables are (1) the strength of the feedback of expectations ($T\#$) and (2) the timing and size of an announced change to γ ($A\#$). Treatments are given in Table 1.

Table 1: Experimental Treatments

Feedback Treatments	Announcements Treatments
T1: $\beta = -0.9$	A1: $\gamma = 60$ for $t = 1, ..., 49$ and $\gamma = 90$ for $t = 50$
T2: $\beta = -2$	A2: $\gamma = 60$ for $t = 1, ..., 19$ and $\gamma = 90$ for $t = 20, ..., 50$
T3: $\beta = 0.5$	A3: $\gamma = 60$ for $t = 1, ..., 19$, $\gamma = 90$ for $t = 20, ..., 44$, and $\gamma = 45$ for $t = 45, ..., 50$.

Using the 3×3 design, we investigate the following hypotheses, which are based on our theoretical results and simulations.

Hypothesis 1 (Stability): *Treatments with $\beta < -1$ result in slower rates of convergence or even non-convergence compared to treatments with $|\beta| < 1$.*

When $|\beta| < 1$, Theorems 1 - 3 imply asymptotic stability of the REE for any specification of the unified dynamic. In addition, the simulations of Section A2 suggest rapid and possibly oscillatory convergence in the T1 treatment and monotonic convergence in T3 treatments. In T2 treatments, where $\beta = -2.0$, results from Theorem 3 and from simulations suggest that asymptotic coordination on the REE is challenging under unified dynamics.[20]

Hypothesis 2 (Level-k Reasoning): *Participant's predictions in announcement periods in treatments A1 - A3 follow level - k deductions for all treatments.*

The announcement treatments, A1 - A3, allow us to precisely identify if agents form high order beliefs following level-k deductions because the rounds played before an announcement's implementation provide an anchor for level-0 forecasts. In other words, the fact everyone observing market price dynamics, as well as its near convergence to the REE, provides an obvious level-0 forecast from which to make level-k deductions. Figure 1 illustrates k-level heterogeneity of individual forecasts; consequently, forecasts should diverge from each other after the announcement. This allows us to precisely characterize whether individual forecasts coincide with the unified model. Importantly, we do not impose, or inform subjects of, a level-0 forecast so coordination on a shared adaptive level-0 forecast is an integral part of the hypothesis.

Hypothesis 3 (Replicator Dynamics): *In response to losses, some participants revise their k-level (up or down) to the current optimal predictor.*

Under unified dynamics, agents who revise their depth of reasoning choose the optimal predictor based on the last period's price. For some agents, this may result in a reduction in reasoning depth. This reduction may be particularly acute following structural change: see Figure 1.

Hypothesis 4 (Level-k Dynamics): *The average depth of reasoning is increasing over time for treatments T1 and T3, during periods when the structure is unchanged. The depth does not increase in the T2 treatments.*

This hypothesis derives its intuition from results on replicator dynamics – see Theorem 2. In particular, if $|\beta| < 1$ (i.e. T1 and T3) then $\hat{k} \to \infty$ whereas if $\beta < -1$ (i.e. T2) then $\hat{k}$ is bounded.

[20] Bao and Duffy (2016) note that the T2 treatment does not satisfy eductive stability, which implies that agents should be unable to coordinate on the REE price. Separately, and as also noted by Bao and Duffy (2016), when the number of participants in a market is finite, the eductive stability condition is relaxed to $-N/(N-1) < \beta < 1$: see Gaballo (2013). Therefore, the appropriate condition for our experiment is $-6/5 < \beta < 1$. The T1×A1 and T2×A1 treatments also serve as a replication exercise for Bao and Duffy (2016).

Figure 2: Screenshot of experimental market GUI

Market Game

What is your prediction?

Submit

Period: **10 of 50**

Period	Your Prediction	Market Price	Forecast Error	Your Earnings
1	20.0	7.59	-12.41	$0.00
2	25.0	10.15	-14.85	$0.00
3	23.0	18.88	-4.12	$0.00
4	19.0	22.74	3.74	$0.08
5	20.0	19.72	-0.28	$0.50
6	18.5	20.95	2.45	$0.32
7	22.0	18.54	-3.46	$0.14
8	21.0	20.2	-0.8	$0.48
9	21.2	22.81	1.61	$0.42
10				

Total Earnings: **$1.94**

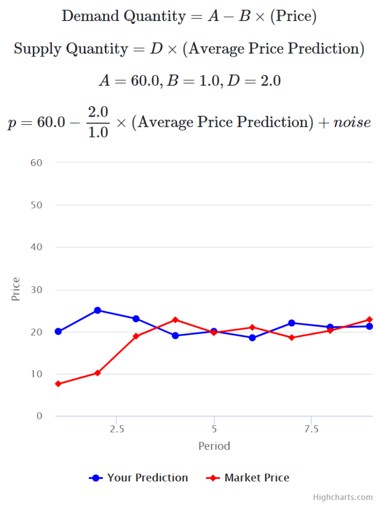

Finally, we note that the four hypotheses, if true, provide evidence against simple alternative models. Standard heuristic switching models, for example, are ruled out by hypothesis 2. Fixed level-k models are ruled out by hypotheses 3 and 4. Purely adaptive dynamics ($\delta_r = 0$) is ruled out by all four hypotheses. Confirmation of the four hypotheses is both evidence for the unified model and against the individual nested alternatives.

5.1 Experiment description

The experiment used a computer based market programmed in oTree. The market setup follows Bao and Duffy (2016) with additions that accommodate our novel elements. Laboratory participants were randomly assigned to groups of six subjects to form markets. Participants were told that they are acting as expert advisers to firms that produce widgets, and were led through a tutorial that describes the market environment including the exact demand and supply equations that govern the price. Participants were informed that the price depends on the average expected price of all advisers in the market and that prices are subject to small white noise shocks.

Participants were given slightly different stories about the market environment in the positive (T3) and the negative (T1 and T2) feedback cases. In the latter, participants were told that the market follows the normal cobweb setup of perfect competition among firms that face convex costs of production of a non-storable good. In the former, participants are told that the widget is a Veblen good with upward sloping demand. In each case, the type of feedback in the market is explained in detail with the paper instructions given to participants containing a version of following text: *"KEY POINT: The market has positive feedback. Therefore, if the average price forecast is high, then the market price will be high. And, if the average price forecast is low, then the market price will be low."* The negative case is stated similarly.[21]

Figure 2 shows the graphic user interface (GUI) that participants interacted with during the experiment. The market information is shown in the top right corner of the screen. A time series plot of the price and the participant's predictions is provided on the bottom right. A table with past prices, predictions, forecast errors, and earnings is provided on the left-hand-side of the screen.

The payoff function for the participant's predictions is

$$\text{payment}_t = 0.50 - 0.03\left(p_t - E_{t-1}p_t\right)^2$$

where p_t is the actual market price in the round, $E_{t-1}p_t$ is their prediction for the price in round t, and 0.50 and 0.03 are measured in cents.[22]

Announcements for the changes in γ were introduced using a pop-up box. The pop-up box described the change in parameters and participants were required to close the box before they could continue. The announcement would also appear, highlighted in red, across the top of the screen in the announcement period. The information in the top right corner of the GUI would also reflect the change.[23]

Each participant played 50 rounds. There was no time limit for each round. Afterwards, participants were asked which strategy they employed, which strategy they believed others employed, and the information they found most useful.

6 Experimental Results

In total, 372 individuals participated in 62 experimental markets: see Appendix for summary statistics. The first two columns of Figure 3 provide an overview of the experimental results from the T1×A3, T2×A3, and T3×A3 treatments. These three treatments illustrate the most novel features of our experimental results and we return to this figure multiple times in the subsequent section. The third column shows the unified model's fit to the aggregate experimental data. The last column shows the best fit of the nested models within the unified

[21] We checked for comprehension of the market environment: see Appendix.

[22] Negative quantities receive zero cents: see Appendix for further details on market payments.

[23] A minimum price of 0 and a maximum price of 500 were enforced. These bounds were not stated to participants, but if violated a pop-up box would appear informing them of the bound.

dynamics when either adaptive learning, the replicator, or level-k deductions is omitted. The specific details of these last two columns are discussed in Section 6.5 but shown here for ease of comparison with the actual data and to illustrate the necessity of all elements of the unified dynamics.

6.1 Convergence Results

The second column of Figure 3 illustrates general convergence properties found across the T treatments. T1 and T3 treatments quickly converge a few periods after the experiment begins. Markets destabilize following announcements, but quickly re-converge within a few periods. T2 treatments are much more volatile: convergence takes much longer and individual forecasts continue to vary widely even once the market price is close to the steady state.

To quantify the speed of convergence, we make use of the experimental design where announcements destabilize the market and set off a new period of convergence. This roughly doubles our sample to 111 distinct market periods to study. We measure convergence using three different metrics. We discuss the *ratio metric* here and refer the reader to the Appendix for discussion of metrics based on mean price discrepancy and mean earnings.

Define a round to be converged when the price is within ± 3 of the steady-state price. The ratio metric applied to an interval of consecutive rounds is defined to be the proportion of converged rounds. Columns 2 - 4 of Table 2 show the ratio metrics the three feedback treatments, where we look at three intervals over the first 19 rounds for all treatments and the comparable intervals for rounds 21 through 38 for treatments with an announcement in period 20. We say that a collection of consecutive rounds has converged if the ratio metric is at least 0.85. Bolded values in Columns 2 - 4 of the table indicate failure to converge. By this metric, none of the feedback treatments (T1, T2, and T3) show convergence within the first five periods of the experiment or within five periods after the first announcement. Convergence is achieved though for T1 and T3 treatments over rounds 6 to 10, rounds 26 to 30, and overall for the full intervals. For the T2 treatments, the 85% threshold is never reached. The other convergence metrics provide qualitatively similar results: see Appendix.

Figure 3: Comparing the unified model to experimental data

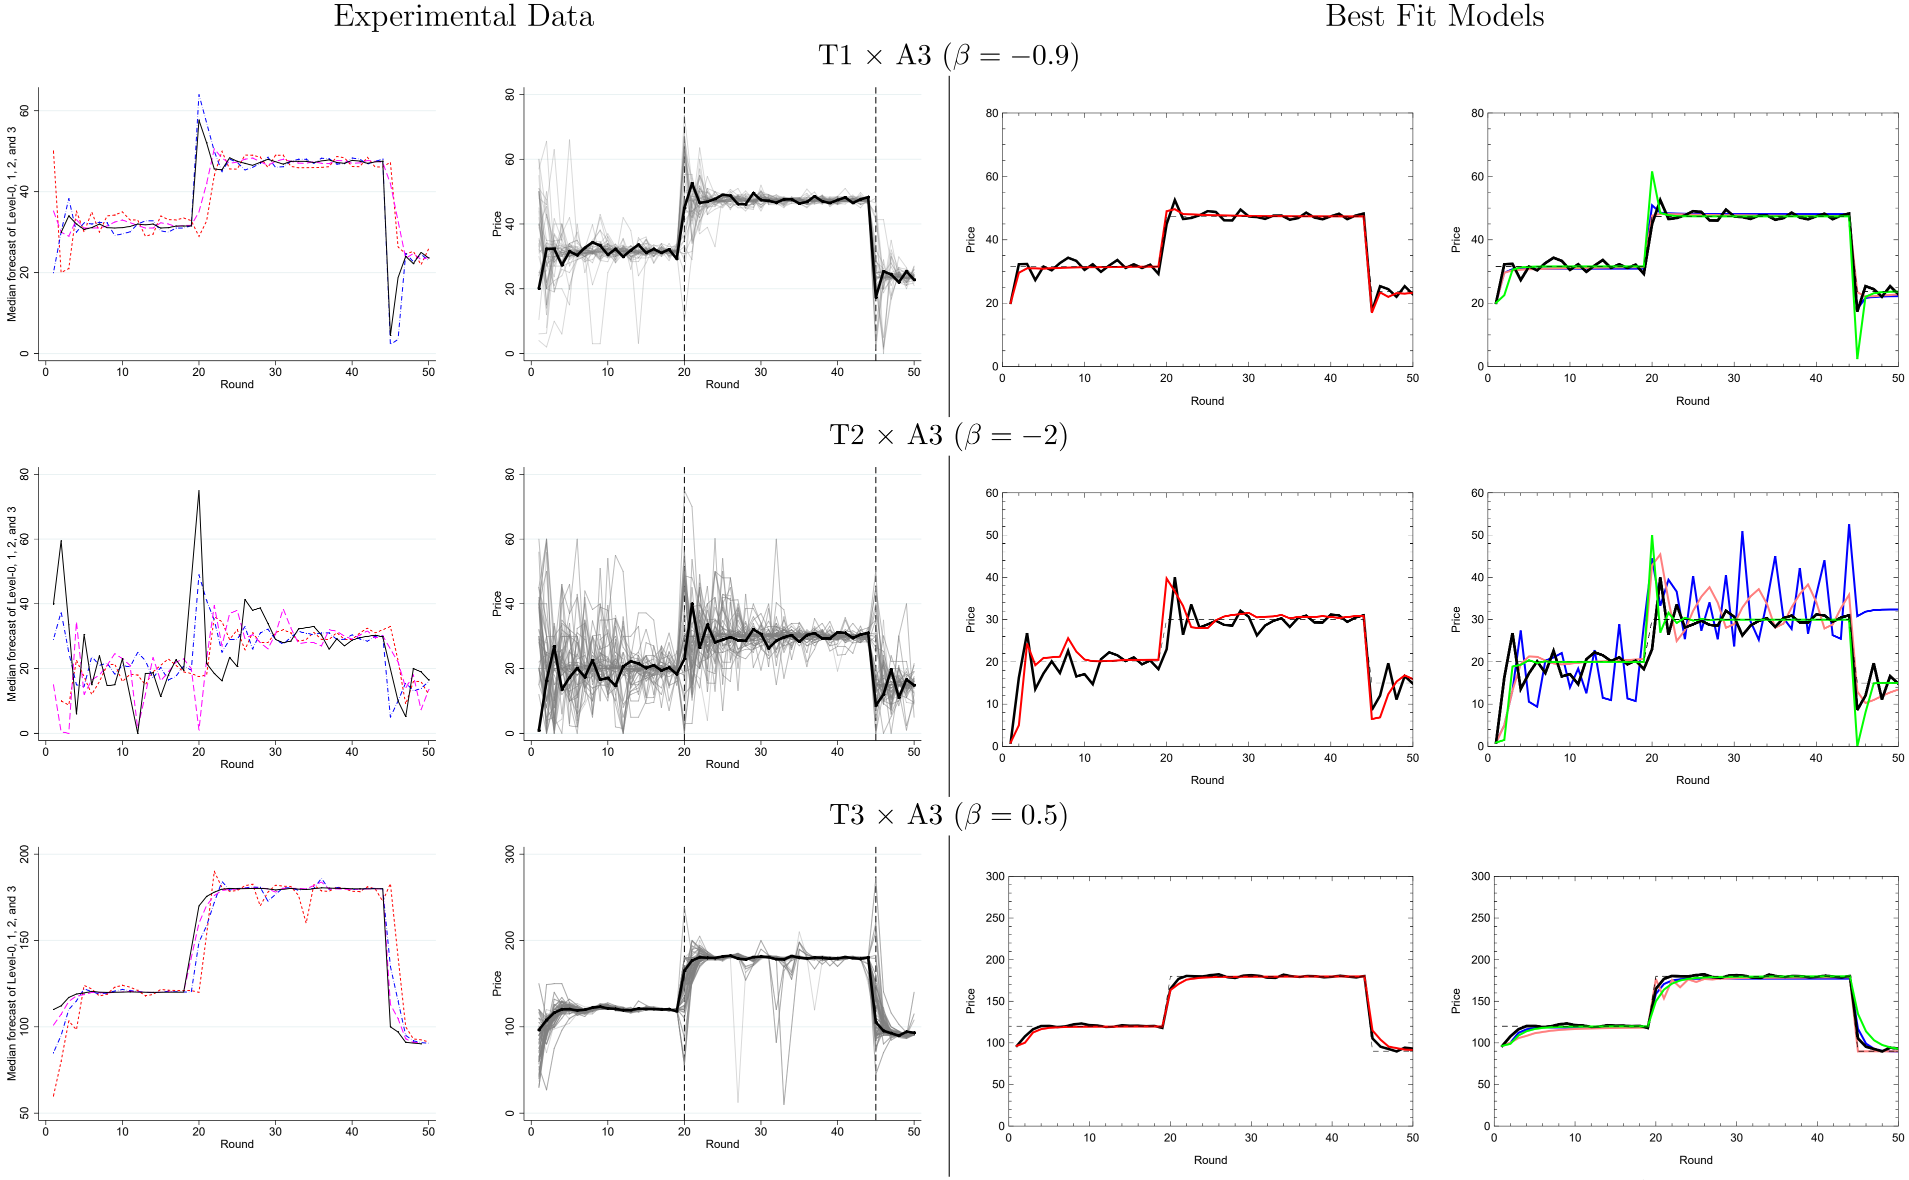

Notes: Survey participants' forecasts are classified as Level-0, 1, 2, 3, or consistent with the REE forecast by comparing to the model implied forecasts. The median forecasts, $E_{t-1}y_t^k$, for $k = 0, 1, 2, 3$ are distinguished by plot-style: *red dotted, blue dash-dot, magenta dash* and *black solid*, respectively. The second column shows average market prices observed (solid black) laid over all individual forecasts. The third column is a fitted unified model with simulated paths initialized to the experimental distribution. The fourth column show best fit alternatives with $\alpha = 0$ (blue), $\phi = 0$ (pink), and adaptive learning only (green).

6.2 Level-k results

A novel feature of our experimental design relative to other level-k studies is that there are many rounds of play before an announcement round. These rounds of play act as a natural reference point to coordinate level-k deductions around a shared level-0 forecast. From this shared level-0 forecast, it is straightforward to predict what types of forecasts we should observe in announcement rounds. In addition, the very first round of play provides a check on this logic. In the first round, there is no shared history to draw upon and no natural level-0 forecast, but is an announcement. Comparing participants' forecasts in round one to those in subsequent announcement periods provides a check for whether participants are coordinating around an adaptive level-0 forecast.

To investigate the degree to which laboratory participants' forecasts follow level-k deductions, we proceed by constructing the implied level-0, 1, 2, 3, and REE forecasts for each experimental market and compare these forecasts to the actual forecasts that laboratory participants submitted. Specifically, we define the level-0 forecast as the average of the two most recent prices.[24] Using this level-0 forecast for each market, we then construct the implied level-1, 2, 3, and the REE forecasts. Then, we calculate the absolute difference between a subject's forecasts in each round and each of the model implied forecasts. We classify each forecast as either level-1, 2, 3, or the REE according to which has the smallest observed difference. Conflicts in classification, if they arise, are resolved by assigning to the lowest level of reasoning. For the first round, when there is no past history of prices, we use the price from the example on the instructions for the T1 and T2 treatments as the level-0 forecast. The modal forecast given by participants in these treatments is close to this value despite no theoretical reason for why people should choose it. For the T3 treatment, we choose the modal forecast observed in the experimental data in round one as the level-0 forecast.

We stop our classification of types at level-3 deductions because higher levels of deduction become hard to distinguish from the REE forecast in the T1 and T3 treatments, and from one another in the T2 treatments in certain settings. We find that approximately 40% of subjects' forecasts that we classified as the REE forecast in a round submit exactly the REE forecast. The remainder are within the ± 3 of it. Therefore, the REE forecast designation likely includes some higher

[24]The results are robust to reasonable changes in the definition of level-0 forecast. In the online Appendix we reproduce all of our results under the five alternative level-0 assumptions including three constant gain learning specifications and find qualitatively similar results. We also explore one market in detail in the online Appendix, which illustrates further how the classification works in practice.

Table 2: Convergence of price to REE in experimental markets

Rounds	Ratio Metric (Converged/Total)			Mean $\lvert p_t - \bar{p} \rvert = \mu$ $H_0 : \mu \leq 3$ $H_a : \mu > 3$			Mean Earning $= \mu$ $H_0 : \mu \geq 0.40$ $H_a : \mu < 0.40$		
A1 - A3	T1 ($\beta = -0.9$)	T2 ($\beta = -2$)	T3 ($\beta = 0.5$)	T1 ($\beta = -0.9$)	T2 ($\beta = -2$)	T3 ($\beta = 0.5$)	T1 ($\beta = -0.9$)	T2 ($\beta = -2$)	T3 ($\beta = 0.5$)
$[2, 5]$	**0.76**	**0.26**	**0.48**	2.22	**9.21**	**5.04**	**0.24**	**0.09**	**0.20**
	(61/80)	(23/88)	(38/80)	[-1.81]	[7.09]	[3.55]	[-9.17]	[-26.26]	[-13.77]
$[6, 10]$	0.97	**0.48**	0.88	1.41	**5.32**	2.08	**0.34**	**0.17**	**0.36**
	(97/100)	(53/110)	(88/100)	[-18.47]	[4.19]	[-2.39]	[-5.21]	[-16.71]	[-4.36]
$[11, 19]$	0.96	**0.73**	0.94	1.18	2.93	1.99	0.41	**0.27**	0.42
	(173/180)	(144/198)	(170/180)	[-9.41]	[-0.26]	[-2.03]	[1.11]	[-10.24]	[2.66]
A2 - A3									
$[21, 25]$	**0.79**	**0.35**	**0.74**	1.86	**6.16**	2.50	**0.25**	**0.12**	**0.28**
	(44/56)	(21/60)	(59/80)	[-3.97]	[4.72]	[-1.16]	[-6.76]	[-14.77]	[-5.89]
$[26, 30]$	1.00	**0.64**	0.94	1.48	**4.22**	1.76	**0.37**	**0.23**	**0.36**
	(70/70)	(48/75)	(94/100)	[-24.21]	[1.61]	[-5.99]	[-3.66]	[-8.73]	[-3.54]
$[31, 38]$	1.00	**0.84**	0.89	0.55	2.52	1.29	0.46	**0.34**	0.42
	(126/126)	(113/135)	(160/180)	[-69.52]	[-1.08]	[-7.46]	[19.53]	[-4.59]	[1.82]
All									
$[2, 19]$	0.92	**0.56**	**0.82**	1.48	**4.99**	2.69	**0.35**	**0.20**	**0.35**
	(331/360)	(220/396)	(296/360)	[-10.91]	[6.48]	[-1.01]	[-6.52]	[-22.85]	[-6.26]
$[21, 38]$	0.95	**0.67**	0.87	1.10	**3.80**	1.70	0.39	**0.26**	**0.37**
	(240/252)	(182/270)	(313/360)	[-24.71]	[2.29]	[-8.05]	[-1.45]	[-12.77]	[-3.39]
Difference	-0.03	-0.12	-0.05	0.38	1.19	0.99	-0.04	-0.06	-0.02
	[-1.67]	[-3.12]	[-1.76]	[2.36]	[2.55]	[2.87]	[-3.41]	[-4.13]	[-1.76]

Bolded values do not meet our criteria for market convergence.

Notes: The table reports three measures of market convergence. Columns 2-4 report the number of rounds where we observe the market price is within ±3 of the REE price. Columns 5-7 report the mean difference between the market price in a round relative to the REE price for the indicated interval of rounds. Columns 8-10 report the mean earning by participants per round over the indicated interval. The maximum earnings in a round is $0.50. We pool A1 (round 50 announcement) and A2 (round 20 announcement) results because both experiments feature a single and identical announcement.

levels of deductions as well.

The upper-left part of Table 3 summarizes the proportion of individuals classified as level-k (for $k = 0, 1, 2, 3$) or REE, for each of the announcement rounds using this ±3 cutoff. The data from all treatments are pooled. The ranges in square brackets show the classification proportions associated with a ±1.5 and ±4.5 cutoff, respectively. Overall, using the ±3 cutoff, we find about half of participants follow a level-k forecast or choose the REE in round one. This number rises to approximately two-thirds for the second and third announcements.

The right side of Table 3 provides a logical check on our classifications. It is natural to think that higher levels of deduction require greater cognitive resources: a person who makes a level-0 forecast might not spend as much time formulating a forecast as someone who makes a level-3 forecast. If our classifications are identifying people who are making level-k deductions we should find a correlation between the time spent deliberating and the depth of reasoning that we identify.

To investigate this, we estimate the following regression model:

$$d_{i,r} = \alpha_i + \omega_r + \sum_k \beta_k I(k)_{i,r} + \sum_k \gamma_k \big(I(k)_{i,r} \times I(Ann)_r\big) + \epsilon_{i,r}, \qquad (7)$$

Table 3: Classifying participant's forecasts as Level-k

Within ±3 of Level-k in announcement rounds	1	20/50	45
Total Classified	47.3%	64.4%	66.0%
	[33.8% , 56.9%]	[52.6% , 71.6%]	[48.1% , 70.5%]
Level-0	14.8%	6.6%	5.1%
	[11.0% , 15.1%]	[4.31% , 8.05%]	[4.49% , 7.05%]
Level-1	7.3%	24.1%	14.1%
	[6.45% , 8.60%]	[20.7% , 26.7%]	[12.2% , 14.1%]
Level-2	6.5%	5.5%	3.8%
	[1.88% , 6.45%]	[4.60% , 5.75%]	[1.92% , 3.85%]
Level-3	3.2%	3.4%	4.5%
	[1.11% , 11.3%]	[2.87% , 4.02%]	[3.85% , 5.13%]
REE	15.6%	24.7%	38%
	[13.4% , 15.6%]	[20.1% , 27.0%]	[25.6% , 40.4%]
N	372	348	156

Hypothesis tests of deliberation time regressions	
H_0 : Level-0 - Level-3 = 0	F(1, 61) =1.87
H_0 : (Level-0 x Ann) - (Level-3 x Ann) = 0	F(1, 61) =**4.59**

Differences in deliberation time (seconds) Variable	(1)	(2)
Level 0	**-5.84**	**-1.31**
	(0.854)	(0.565)
Level-1	**-4.89**	-0.90
	(0.950)	(0.701)
Level-2	**-3.96**	-1.13
	(1.210)	(0.851)
Level-3	**-3.82**	0.21
	(1.370)	(1.113)
Level-0 x Ann	**45.25**	2.44
	(8.763)	(6.021)
Level-1 x Ann	**43.12**	**12.25**
	(4.712)	(4.554)
Level-2 x Ann	**59.58**	12.85
	(8.806)	(8.390)
Level-3 x Ann	**62.83**	**22.31**
	(11.80)	(8.392)
Cons	**39.5**	**112.68**
	(0.457)	(4.206)
Individual FE	yes	yes
Round FE	no	yes
R-squared	0.027	0.253
N	18,367	18,367

Notes: The top left panel reports the proportion of participant's forecasts that fall within ±3 of a Level-k forecast. Proportions for cutoffs of ±1.5 and ±4.5 are shown in brackets. The right panel reports the regression results of identified Level-k individual's deliberation time in all periods and in announcement periods. Standard errors are clustered at the market level and reported in parenthesis below the point estimates. Bolded values indicate statistical significance at the ten percent level. The bottom left panel reports the hypothesis tests for the equality of regression coefficients for regression specification (2). We pool A1 (round 50 announcement) and A2 (round 20 announcement) results because both experiments feature a single and identical announcement.

where $d_{i,r}$ is the time spent deliberating by person i in round r, α_i is an individual fixed effect that controls also for treatment and market, ω_r is a round fixed effect since typically less time is spent deliberating in later rounds, $I(k)_{i,r}$ is an indicator identifying whether person i is classified as choosing a level-k forecast in round r; and $I(Ann)_r$ is an indicator identifying whether an announcement is made in round r. Standard errors are clustered at the market level. The coefficients β_k and γ_k estimate the difference in deliberation time, overall and in announcement rounds respectively, for those identified as level-k for $k = 0, 1, 2, 3$, relative to those whom we identify as choosing the REE forecast or we fail to classify.

The regression results confirm our hypothesis. We find that those whom we identify as level-0 spend the least amount of time deliberating on their forecast overall, and in announcement rounds. Those identified as level-3 spend the most amount of time among the classified types in all rounds, and in announcement rounds, with the difference between deliberation times of level-0 and level-3 participants statistically different at standard significance levels.

Figure 4 shows histograms of individual forecasts in round one and in each announcement round for each feedback treatment. The gray bars show the model implied level-k forecasts with ±3 band. The T2 round 20/50 predictions provides the clearest level-k deductions because the large negative feedback ($\beta = -2$) makes each level-k prediction very distinct. While there is not strong evidence for level-k reasoning in the first period,[25] this changes once participants have played multiple rounds and an announcement occurs. For these announcement rounds, Figure 4 and Table 3 show a majority of participants playing level-k or the high level-k/REE forecasts.

The exit surveys support this interpretation. On average, participants claim that the equations and a forecast of average expectations were more important for them than for other participants. Further, participants claim that past prices were more important for others' forecasts than their own. In the interest of space, the survey results are discussed in the online Appendix.

6.3 Negative Expectational Feedback in Macroeconomic Modeling

Our experimental results also shed light on an element of level-k reasoning that recently has been called into question in the literature. García-Schmidt and Woodford (2019) and Angeletos and Sastry (2021) put forward models of bounded rationality that modify level-k reasoning to rule out oscillating deductions when there is negative expectational feedback. Angeletos and Sastry (2021) writes, "We are not aware of any experimental evidence of this oscillatory pattern. We suspect that it is an unintended "bug" of a solution concept."

In our experiments, we provide evidence that oscillating deductions occur. Figure 5 is based on a T2×A2 treatment. The NW panel shows market price in bold and players' forecasts in light gray; the N panel shows market price in bold and the model implied level k forecasts (note truncated time-span around announcement). The NE panel shows players' forecasts with dots colored to identify level-k classification. Remaining panels provide model implied level-0 forecasts in dashed red and player forecasts colored to identify level-k classification.

From Figure 5 is evident that some individual participants make oscillating deductions over time. Obvious examples include participants #1 and #4. Indeed, four out of six of the participants clearly make oscillating deductions with forecasts above and below the REE after the announcement occurs in round 20. The oscillations occur despite the experience of the price not oscillating for many periods prior to the announcement. This experience of tranquility combined with how close many of the forecasts are to level-k deductions suggests that participants contemplated oscillations consistent with classic level-k reasoning. Moreover, they took action with money at stake which was consistent with such deductions.

[25] This lack of evidence may reflect difficulty in establishing level-0 forecasts in round one. Also, some participants do not appear to understand the game's structure in the first period.

Figure 4: Laboratory subjects' forecasts in announcement rounds

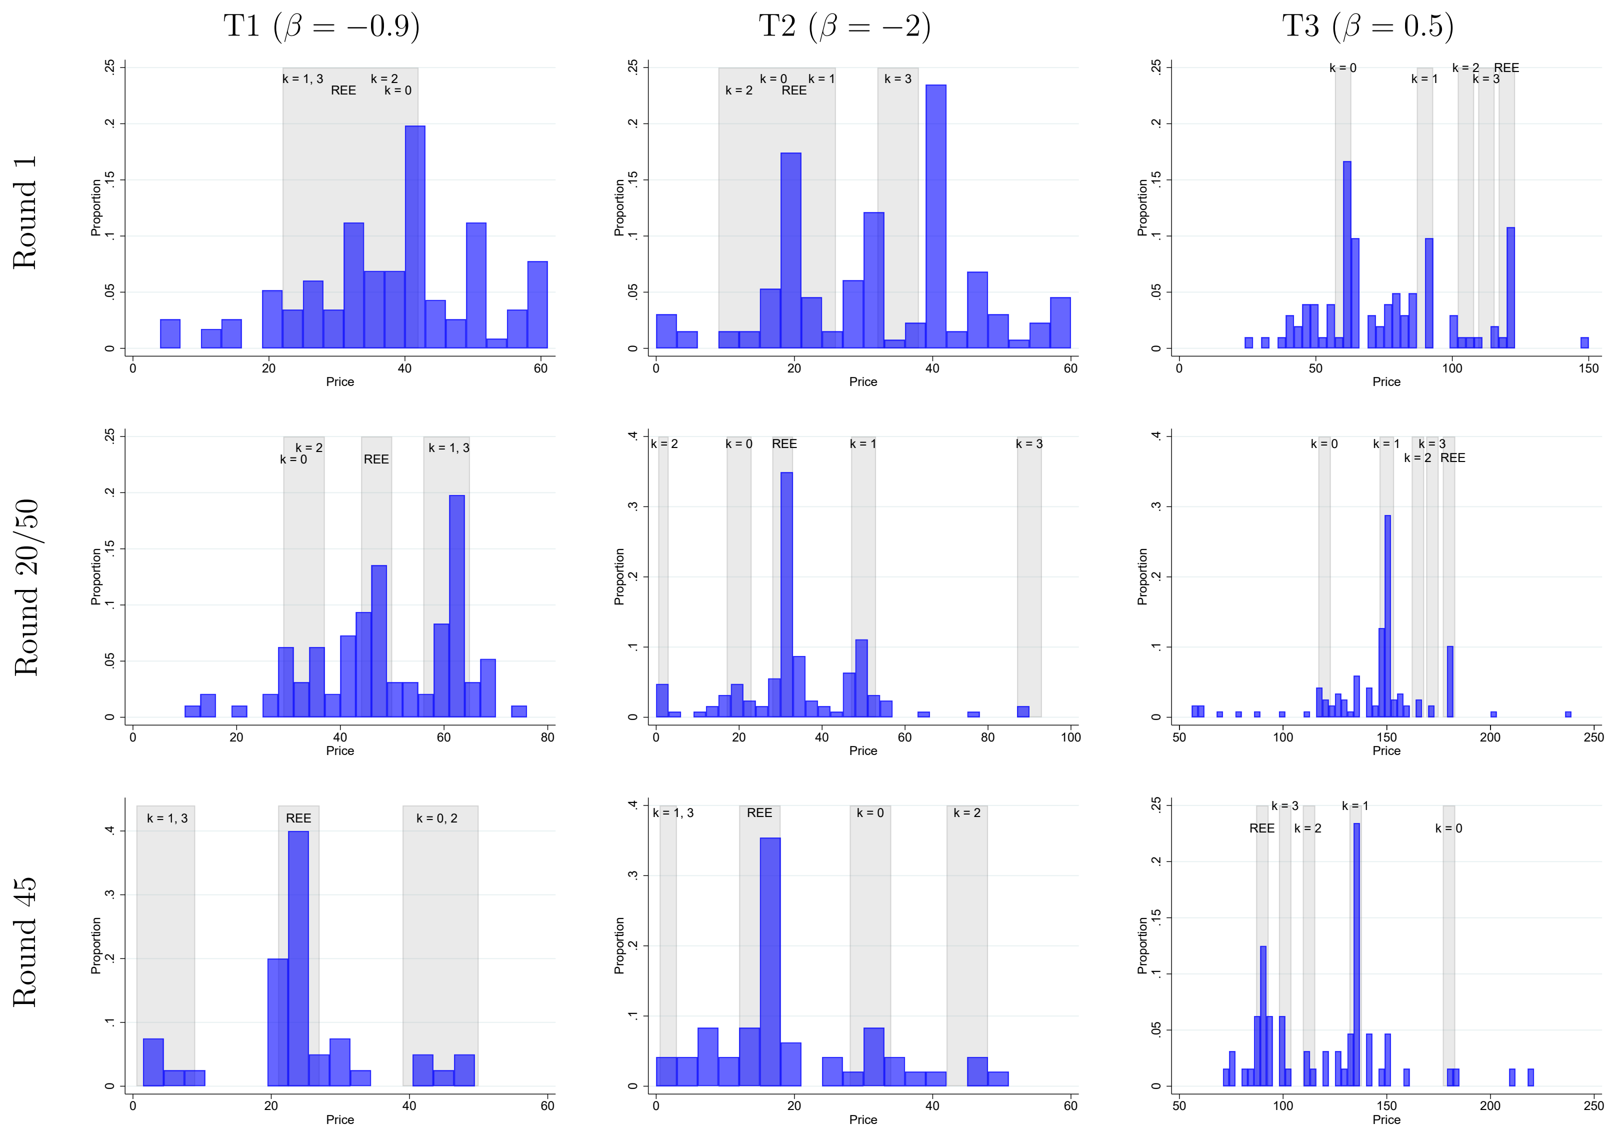

Notes: Histograms of the subject's forecasts in response to an announced structural change. The shaded regions correspond to our classifications of level-0, 1, 2, 3, and the REE forecasts reported in Table 3, which is ±3 of the model implied Level-k forecast. The width of each bin for the experimental data is 3. The level-0 shaded bar includes the previous steady state for prices prior to the announcement in round 20/50 and round 45 cases. We pool A1 (round 50 announcement) and A2 (round 20 announcement) results because both experiments feature a single and identical announcement.

Figure 5: Oscillating Deductions: Individual forecasts from experimental market with treatment T2 ($\beta = -2$)

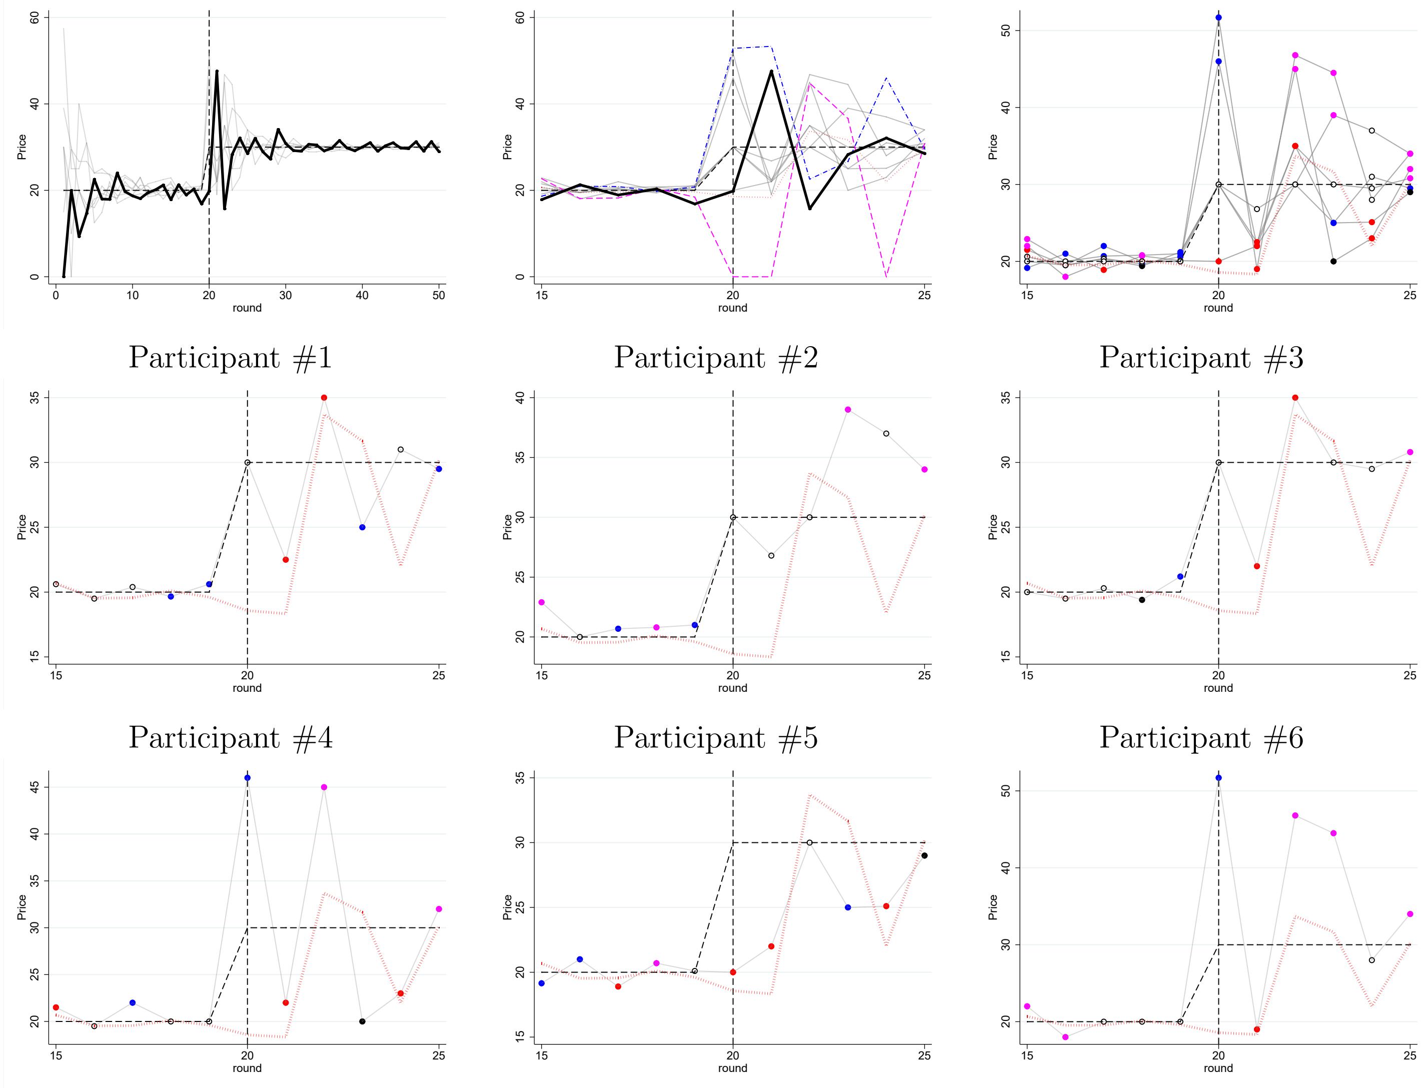

Notes: The first plot shows all individual forecasts and the market price from a single market. The second plot shows the model implied level-k forecasts and the data in a window around the announcement. The remaining figures classify each of the forecasts as a level-k type, which is indicated by the color of the dot. Forecasts that are classified as level-0 are shown in red, level-1 in blue, level-2 in magenta, level-3 in black, and REE as a black circle with a white interior. The dotted red line shows the path of the level-0 forecast from which all level-k deductions are derived. The perfect foresight equilibrium is indicated by the dashed line.

6.4 Revisions to the depth of reasoning

The replicator employs three key assumptions. First, in any given period and for any level k, some participants maintain their depth of reasoning. Second the proportion k-level reasoners who revise their depth of reasoning is increasing in the size of the most recent forecast error. Third, participants who revise their depth of reasoning choose the level that would have been optimal last period.

To test the three features of the replicator dynamic, we make use of the announcements in the A2 and A3 treatments. The announcement rounds provide a clear intervention from which to identify level-k deductions. The structural change leads to large forecast errors for many participants, and also provides distinct counterfactual level-k predictions. These can be used to identify revisions to the depth of reasoning in the following period.[26]

Table 4 reports the results for the first and second announcements across all treatments. The first column shows the proportion of individuals who, conditional on changing strategies in the period following the announcement, are classified as selecting the best counterfactual strategy from the previous period. We note that, in accordance with the unified model, we often observe a proportion of agents lowering their reasoning levels in the period following announcements.[27] The second column reports the proportion of participants whom we identify as not changing their strategy. The remaining columns report the difference in mean absolute forecast errors experienced by changers and non-changers, and the deliberation time when selecting their new forecast.

We find evidence consistent with our replicator assumption for all three key aspects. First, we document that a proportion of subjects do not update their strategy following the announcement period. Second, the subjects changing strategy experienced larger forecast errors (on average) and spent more time deliberating. The T3 x A2/A3 treatment does not fully follow the predicted pattern: changers make larger forecast errors, but spend less (but not statistically significantly less) time deliberating. Finally, a significant proportion of those who do change strategy choose the previous period's optimal k-level strategy.

The unified model also predicts that when $|\beta| < 1$ we should see increasing depth of reasoning over time during periods when the market structure is constant. We can test this prediction by looking at the distribution of strategies that are played across the same subjects in the A3 treatments with two announcements. The unified model predicts that over time more people will select higher level-k forecasts for the T1 and T3 treatments, but not for the T2 treatments.

[26] The classification of forecasts is restricted to level-0, 1, 2, 3, and REE, and is given by the level-k strategy nearest in mean squared error to the submitted forecast.

[27] For the T1 treatment, for example, the proportion is 20.2%.

Table 4: Revisions and loss

	Proportion of changers Between rounds 20 & 21		Ave. abs. prediction error Round 20			Ave. deliberation time (sec) Round 21		
Treatment	Revise opt.	No Change	Change	No change	Difference	Change	No change	Difference
T1 x A2/A3	**0.40**	0.38	17.82	8.44	**9.37**	56.8	50.5	6.38
	[2.35]	(32/84)			[4.53]			[0.72]
T2 x A2/A3	**0.35**	0.49	23.07	14.59	**8.48**	64.3	54.6	9.76
	[1.41]	(44/90)			[3.04]			[1.11]
T3 x A2/A3	**0.55**	0.31	29.21	11.66	**17.31**	34.8	39.8	-4.99
	[5.74]	(37/119)			[5.69]			[-0.82]
	Between rounds 45 & 46		Round 45			Round 46		
T1 x A3	**0.68**	0.55	24.43	3.13	**21.3**	42.1	28.6	**13.5**
	[4.02]	(23/42)			[7.60]			[1.57]
T2 x A3	0.24	0.40	18.83	6.95	**11.88**	31.6	30.7	0.85
	[-0.10]	(19/48)			[4.22]			[0.14]
T3 x A3	**0.41**	0.26	30.15	28.5	1.66	26.0	19.3	**6.71**
	[2.35]	(17/66)			[0.19]			[2.24]

Notes: "Revise opt." is the proportion of people who, conditioning on changing their strategy in period 21(46), changed their strategy to the best counterfactual strategy out of level-0, 1, 2, 3, or the REE in their market, where best is defined as what forecast would have been best in round 20(45). Z-scores for the test of the null hypothesis that subjects switched to one of the five strategies at random are reported in brackets. The next column reports the proportion of participants who we classify as not changing their strategy either between rounds 20 and 21 or between rounds 45 and 46 following announcements in either round 20 or 45, respectively. Counts appear in parentheses below. The remaining columns report the difference in average absolute prediction errors and average deliberation time for subjects classified as changing versus not changing with two-sample t-test statistics reported in brackets. Bolded values represent statistical significance at the ten percent level.

Figure 6 shows the distribution of forecasts for levels-0 to 3 and REE. We can see for the T1 and T3 cases that the distribution shifts to the right. More subjects choose higher-level forecasts, or are consistent with the REE forecast in the second announcement than in the first. We find that a Kolmogorov-Smirnov equality of distributions test rejects the null of equality at the 5% level for the T1 and T3 treatments. The T2 treatments, however, shows a different result. For T2 treatments, we observe a bifurcation in which subjects either choose a low levels of reasoning or they jump to the REE.

6.5 Quantitative Evaluation

So far, we have focused on individual-level behavior. In this section, we use aggregate price data to compare the fit of the unified model to the fit of simpler alternatives models: a fixed level-k model, a replicator-only model, a pure adaptive learning model, and REE.[28]

We use our classification of level-k types in period one to initialize the models in each market. The fixed level-k model allows for the level-0 forecast to evolve over time with the proportion of agents using different level-k types fixed to the initial values. The replicator-only model assumes a fixed level-0 forecast but allows for the choice of level-k forecasts to vary over time. The adaptive learning

[28] For each model (except REE), we computer the forecast parameters minimizing the squared error between the simulated data and the experimental data.

Figure 6: Increasing or decreasing depths of reasoning over time

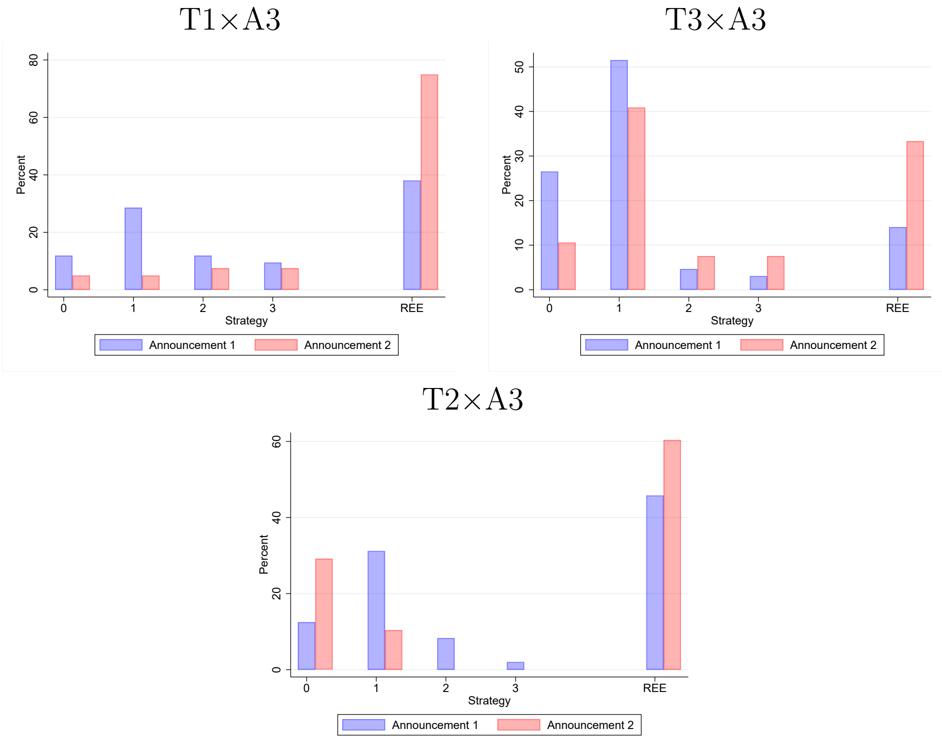

Notes: The distribution of classified types of forecasts observed in the experimental treatments with two announcements in each announcement period.

model shuts down the replicator and assumes all agents use the same level zero forecast, which evolves over time as new data become available.

For each version of the model and for each individual market, we compute the mean-squared error (MSE) measured as average over time of the squared difference between the market price and the price obtained by simulating the model using the fitted values of the associated learning parameters. Table 5 shows the average of the MSEs across the T1×A3, T2×A3, and T3×A3 markets. For given treatments, the average prices across markets, and simulations from each fitted model, are shown in the right two columns of Figure 3. In the interest of space, the individual market outcomes are reported in the online Appendix.

Using the Wilcoxon ranked-sign test, we compare the individual-market MSEs of the unified model to those of the alternative models. For each of the expectational feedback treatments, the median MSE of the unified model is lower than of the adaptive learning model at the 10% level; and for the T1 and T3 treatments it is lower than for the REE at the 5% level.[29]

The unified model also outperforms the fixed level-k model and the replicator-only model across all three treatments at least the 5% level. This result, along

[29]For the T2 treatments, the test fails to reject the null hypothesis for equality compared to the REE. We note that the realized market prices are not induced by participants having rational forecasts: see second column of Figure 3.

Table 5: MSE between experimental data and competing models

Treatment	RE	Unified Model		Fixed Level-k		Replicator only		Adaptive learning	
T1 × A3 ($\beta = -0.9$)	MSE	MSE	Rel. RE	MSE	Rel. RE	MSE	Rel. RE	MSE	Rel. RE
Ave. of All Markets	13.15	5.95	0.45	12.37	0.94	9.80	0.74	22.26	1.69
T2 × A3 ($\beta = -2$)									
Ave. of All Markets	51.82	48.38	0.93	422.71	8.16	70.98	1.37	63.39	1.22
T3 × A3 ($\beta = -0.5$)									
Ave. of All Markets	37.17	19.83	0.53	20.78	0.56	49.44	1.33	50.51	1.36

Notes: Average mean square error (MSE) of five simulated models of aggregate price dynamics compared to experimental market price data. "Rel. RE" reports the MSE of the a model relative to RE MSE, i.e., Model MSE/RE MSE. Individual market MSEs that underlie the averages in this table are shown in Table A13 in the Appendix. Models are fit by doing a grid search over values $\alpha \in [0, 2]$ and $\phi \in [0, 1]$.

with the improvement over pure adaptive learning alone, provides evidence that all three elements – adaptive learning, level-k reasoning, and the replicator – are required to explain the aggregate data.

6.6 Discussion

The experimental evidence displayed in Table 2 provides strong support for Hypothesis 1 (stability): large negative feedback results in slow convergence, or nonconvergence, to the REE price, while convergence is achieved for $|\beta| < 1$. The speed of convergence appears to increase following an announcement treatment. Increases in convergence speed in treatments T1 & T3 align with increases in the depth of reasoning when there are multiple announcements: see Figure 6.

We find strong support for Hypothesis 2 (level-k reasoning). In announcement periods, we observe clear bunching around particular k-level forecasts: see Figure 4. Comparing individual and model-implied forecasts in announcement rounds, we classify between 50% and 70% of subjects, as Level-0, 1, 2, 3, or REE. Our classifications align with deliberation times: level-0 participants spend less time deliberating than level-3.

We find support for Hypothesis 3 (replicator dynamics). Focusing again on announcement periods, we find that some fraction of subjects are classified as using the same depth of reasoning in the announcement period and in the period following the announcement. These subjects on average had lower forecast errors in the announcement period than those subjects who appear to change strategies, and they spent less time deliberating in the next round. In addition, for those we classify as changing their strategy, we find evidence that a high proportion are changing to the best strategy (see Table 4). As predicted by our theory, many of these changes correspond to decreases in the subject's k-level depth of reasoning.

We find evidence for Hypothesis 4 (level-k dynamics) for the T1 and T3 treatments: we observe revisions over time in depth of reasoning. Also, the proportion of high-level forecasts increased (inducing quicker convergence) with second announcement: see Figure 6 and Table 2.[30]

7 Unified Dynamics in the New Keynesian Model

The economic environment that we have studied thus far is univariate and relies on one-step ahead expectations. The microfoundations of most macroeconomic models of the business cycle, however, imply agents must form expectations over multiple future variables. To show that our theoretical and experimental results are useful for understanding these more complicated environments, we employ unified dynamics in the standard New Keynesian model to investigate forward guidance monetary policy.

Level-k reasoning has been proposed as a solution to the *forward guidance puzzle* in New Keynesian models, in which credible promises of future monetary policy are found to be implausibly powerful under RE. Angeletos and Lian (2018) or Farhi and Werning (2019) show the puzzle can be resolved if agents are exogenously assumed to be low-level reasoners. We use our model, which endogenizes reasoning levels, to study this issue. We show that under forward guidance, low-level reasoning organically emerges, and we find that the three mechanisms of unified dynamics interact to resolve the puzzle.

We proceed in three steps. First, we introduce the model and the forward guidance policy problem confronting monetary policymakers. Second, we show how a special case of this environment reduces to a univariate model closely related to the model studied in the previous sections. Finally, we examine the implications of an exogenous shock that puts the economy at the zero lower bound. We find that low level-k reasoning can be an endogenous outcome, which substantially lowers the power of monetary policy promises.

7.1 Forward guidance policy problem

We consider the standard New Keynesian economy as described in Woodford (2003). Under RE, household and firm decisions are approximated by IS and

[30] We observe a bifurcation in the distribution of classified strategies played in the T2×A3 treatments between the two announcement rounds, with more level-0 and REE forecasts played in the second announcement round. With strong negative expectational feedback, REE forecasts cannot be viewed as limiting values of level-k forecasts, therefore the increase in the proportion of players classified as using REE does not formally contradict Hypothesis 4. We speculate that the high proportion of REE forecasts observed in the T2×A3 treatment's second announcement round may be due to negative prices are not allowed. In an announcement round, many high-level forecasts predict either 0 or γ in the T2 treatment. Therefore, a subject's menu of forecasts has a finite number of distinct choices. With finite choices and bounded prices, its plausible that some subjects will engage in sufficient reflection to engender more coordination on the REE, which is in the interior of the price space. This is an interesting avenue for future research.

Phillips curve relationships:

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^n) \quad (8)$$
$$\pi_t = \xi E_t \pi_{t+1} + \kappa x_t, \quad (9)$$

where x_t is the output gap, π_t inflation, i_t is the nominal interest rate, r_t^n is the natural real rate of interest, ξ is the discount factor, σ is the intertemporal elasticity of substitution, and κ is a composite parameter that is determined by the degree of price rigidity in the economy.

The exogenous driver of the economy is a Markov process with states S (stagnation) and N (normal) (known to all agents), which determines the natural rate r_t^n. In the stagnation state, $r_t^n = r_S < 0$ and in the normal state $r_t^n = r_N > 0$. For experiment under consideration we assume that, in period zero the economy unexpectedly enters the stagnation state, and it remains there each period with probability $1 - \delta$.

The policy problem is how should policymakers respond to this unanticipated shock. We assume central bank seeks to minimize the following loss function

$$\min E_0 \left\{ \frac{1}{2} \sum_{t=0}^{\infty} \xi^t \left(\pi_t^2 + \psi_x x_t^2 \right) \right\} \quad (10)$$

subject to (8), (9), and $i_t \geq 0$.

Optimal discretionary policy is to set the nominal rate at zero in the stagnation state. To study forward guidance, we follow Bilbiie (2019) and assume that policy makers engage in a partial commitment strategy: in state S the bank announces that it will continue to hold the interest rate at zero beyond the end of state S, i.e., provide forward guidance. The implementation of this policy involves a probabilistic return to normalcy: after the natural rate returns to its normal value of r_N policy makers continue to hold the interest rate at zero each period with probability $1 - \nu$.

With policy modeled in this fashion, the economy is now driven by a three-state Markov process with states S, F (forward guidance) and N, and with transition matrix

$$P = \begin{pmatrix} 1-\delta & \delta(1-\nu) & \delta\nu \\ 0 & 1-\nu & \nu \\ 0 & 0 & 1 \end{pmatrix}.$$

Forward guidance policy is thus reduced to a single parameter choice: ν.

7.2 Unified dynamics in the New Keynesian Model

To develop unified dynamics in this multivariate environment, we assume that the aggregate expectations operators are the average expectations across agents, and that equations (8) and (9) are taken as the current period best response functions over which the agents do level-k deductions.[31] This is a natural benchmark here in part because it is the standard assumption employed in macroeconomic laboratory experiments that test expectation formation in the New Keynesian model such as in Mokhtarzadeh and Petersen (2021) or Kryvtsov and Petersen (2021).[32]

To illustrate how level-k deductions work in this setting, and to facilitate connections to our previous analysis, we begin by deriving these deductions in the special case $\kappa = 0$ in which prices are fixed so that inflation and inflation expectations are zero. The more general case is derived in the appendix.

Assume that in period $t \geq 0$ the economy is (still) in state S and discretionary monetary policy is pursued (ignoring forward guidance for now). We assume that, in state S, level-0 agents hold forecasts $E_t^0[x_{t+1}|S] = a_{t-1}^x$, and that in state N, level-0 agents hold forecasts $E_t^0[x_{t+1}|N] = 0$. Further, we assume that level-1 agents assume that all agents are level-0. Thus for given level zero expectations a_{t-1}^x of the output gap in state S, level-1 forecasts are obtained as follows:

$$\begin{aligned} E_t^1[x_{t+1}|S] &= (1-\delta)E_{t+1}^0[x_{t+2}|S] + \delta \underbrace{E_{t+1}^0[x_{t+2}|N]}_{0} + \sigma(1-\delta)r_S \\ &= (1-\delta)a_{t-1}^x + \sigma(1-\delta)r_S. \end{aligned}$$

Here, for example, $E_{t+1}^0[x_{t+2}|N]$ is the period $t+1$ forecast of x_{t+2} made by level-0 reasoners, given that the state is N in period $t+1$, with this notation being extended in the obvious way.

Level-2 agents assume that other agents are level-1, thus

$$\begin{aligned} E_t^2[x_{t+1}|S] &= (1-\delta)E_{t+1}^1[x_{t+2}|S] + \delta E_{t+1}^1[x_{t+2}|N] + \sigma(1-\delta)r_S \\ &= (1-\delta)\left((1-\delta)a_{t-1}^x + \sigma(1-\delta)r_S\right) + \sigma(1-\delta)r_S. \end{aligned}$$

Continuing in this way, we can define $E_t^0[x_{t+1}|S] = a_{t-1}^x$, $E_t^1[x_{t+1}|S] = T(a_{t-1}^x|S) \equiv (1-\delta)\sigma r_S + (1-\delta)a_{t-1}^x$, and

$$E_t^k[x_{t+1}|S] = T^k(a_{t-1}^x|S) \equiv T\left(T^{k-1}(a_{t-1}^x|S)\right) \text{ for } k \geq 2.$$

[31] In environments with long-lived agents there are a variety of ways to model decision-making: for an alternative implementation see Preston (2005), and for further discussion see Evans and McGough (2020) and Evans and McGough (2021).

[32] Angeletos and Lian (2018) and Farhi and Werning (2019) develop alternate implementations of k-level deductions in New Keynesian environments.

Combining these definitions with the IS equation and substituting in aggregate beliefs yields the realized value of x in state S as a function of level-0 beliefs, i.e. $x_t = \mathcal{T}(a^x_{t-1}|S)$, where

$$\mathcal{T}(a^x|S) = \sigma r_S \left(1 + \frac{(1-\delta)}{\delta} \sum_{k\geq 0} \left(1-(1-\delta)^k\right)\omega_k\right) + \left(\sum_{k\geq 0}(1-\delta)^k\,\omega_k\right) a^x. \quad (11)$$

We note $\mathcal{T}$ is linear in a^x and that (11) is the analog to equation (2) in Section 2.

In the Appendix, we show that level-k deductions expand to include the forward guidance policy and to the multivariate case in the same way. The T-map that governs level-k deductions in this case is

$$T\begin{pmatrix} a^x_{|S} \\ a^\pi_{|S} \\ a^x_{|F} \\ a^\pi_{|F}\end{pmatrix} = \begin{pmatrix} (1-\delta)\begin{pmatrix} r_S\sigma \\ r_S\kappa\sigma\end{pmatrix} + \delta(1-\nu)\begin{pmatrix} r_N\sigma \\ r_N\kappa\sigma\end{pmatrix} \\ (1-\nu)\begin{pmatrix} r_N\sigma \\ r_N\kappa\sigma\end{pmatrix}\end{pmatrix}$$

$$+\begin{pmatrix} (1-\delta) & \sigma(1-\delta) & \delta(1-\nu) & \sigma\delta(1-\nu) \\ \kappa(1-\delta) & (\xi+\kappa\sigma)(1-\delta) & \kappa\delta(1-\nu) & (\xi+\kappa\sigma)\delta(1-\nu) \\ 0 & 0 & (1-\nu) & \sigma(1-\nu) \\ 0 & 0 & \kappa(1-\nu) & (\xi+\kappa\sigma)(1-\nu)\end{pmatrix}\begin{pmatrix} a^x_{|S} \\ a^\pi_{|S} \\ a^x_{|F} \\ a^\pi_{|F}\end{pmatrix}$$

Turning now to the replicator, denote the state in time t by $z_t \in \{S, F\}$. Next, recall that if, in period $t-1$, the state is either S or F, the forecasts of level zero agents *are not* conditional on the realization of the state in period t, i.e. $E^0_{t-1}[y_t] = a^y_{t-1}, y \in \{x, \pi\}$, regardless of the value of z_t. However, for $k \geq 1$, level-k reasoners understand the economy's structure and incorporate it into their forecast behaviors. Thus, agents using level-k reasoning for $k \geq 1$ make forecasts in period $t-1$ that *are* conditional on the realization of z_t, i.e. $E^k_{t-1}[y_t|z_t]$.

In the New Keynesian model agents make forecasts of both the output gap and inflation. Therefore, agents have two forecast errors to consider when assessing the appropriate level-k strategy to choose. To accommodate this change, we assume that the agents evaluate the following loss function

$$\mathbb{L}^k_t(z_t) = \left|\hat{E}^k_{t-1}[\pi_t|z_t] - \pi_t\right| + \psi_x\left|\hat{E}^k_{t-1}[x_t|z_t] - x_t\right|, \quad (12)$$

where $0 < \psi_x \leq 1$ is the same weight the central bank applies to deviations of inflation and the output gap from target. The state-contingent time t optimal predictor is given by

$$\hat{k}(x_t, \pi_t, z_t) = \min \arg \min_{k \in \mathbb{N}} \mathbb{L}_t^k(z_t), \tag{13}$$

where the left-most "min" is used to break ties just as before.

When $z_t = S, F$, unified dynamics in the NK model are given as

$$\begin{aligned}
x_t &= \sum_{k \geq 0} \omega_{kt} E_t^k x_{t+1} - \sigma(i_t - \sum_{k \geq 0} \omega_{kt} E_t^k \pi_{t+1} - r_t^n) \\
\pi_t &= \xi \sum_{k \geq 0} \omega_{kt} E_t^k \pi_{t+1} + \kappa x_t \\
\omega_{it+1} &= \begin{cases} \omega_{it} + \sum_{j \neq \hat{k}(y_t)} r\left(\mathbb{L}_t^j(z_t)\right) \omega_{jt} & \text{if } i = \hat{k}(x_t, \pi_t, z_t) \\ \left(1 - r\left(\mathbb{L}_t^k(z_t)\right)\right) \omega_{it} & \text{else} \end{cases} \\
a_t^x &= a_{t-1}^x + \phi(x_t - a_{t-1}^x) \\
a_t^\pi &= a_{t-1}^\pi + \phi(\pi_t - a_{t-1}^\pi),
\end{aligned} \tag{14}$$

where the final two equations capture the adaptive dynamics of level-0 reasoners.

7.3 Low-levels of deduction and forward guidance

The crux of the forward guidance puzzle is that general equilibrium effects of anticipated policy are too strong under RE. Level-k reasoning offers an illustration of why this occurs. Consider the univariate NK special case, i.e.ẇhen prices are fixed and $a^\pi = 0$. Level-k forecasts for the output gap are given by $E_t^0 x_{t+1} = a^x$, $E_t^1 x_{t+1} = \sigma(1-\nu) r_N + (1-\nu) a^x$, and

$$E_t^k x_{t+1} = \frac{1 - (1-\nu)^k}{\nu} \sigma r_N + (1-\nu)^k a^x. \tag{15}$$

Since $0 \leq \nu \leq 1$, it straightforward to show that $E_t^k x_{t+1}$ is bounded between 0 and $k\sigma r_N + a^x$. In the RE limit $(k \to \infty)$ we see that $E_t^k x_{t+1} \to \sigma r_N / \nu$, which in unbounded as $\nu \to 0$, i.e. as the forward guidance period is extended. Thus, by equation (8), forward guidance can provide infinite stimulus under RE. However, if the average level of reasoning is low, the power of forward guidance is reduced.

As is evident in this special case, forward guidance policy controls the economy's expectational feedback, with $1-\nu$ playing a role analogous to the parameter β found in the univariate model of Section 3. Therefore, applying the intuition obtained from Proposition 1, we would expect that, with uniform weights and $\kappa = 0$. the bound on the optimal reasoning level is approximately 50% of the highest level in use. Moreover, because a lower-level response is optimal, agents using higher level-k reasoning will tend revise to a lower level-k in a dynamic setting, leading to persistence of low-level reasoning.

Moving to the bivariate NK model, we find that low levels of reasoning can also be persistent in the face the adverse shock coupled with forward guidance. This persistence is the result of competing tensions on the realized forecasts for different levels of reasoning. Specifically, for our calibrations, a negative shock coupled with forward guidance results in forecasts that are non-monotonic in k. To examine this point explicitly, consider inflation forecasts $E_0^k \pi_1$ at the time of the shock and for different values of k. At low level-k, both the shock and policy in this environment have small contemporaneous effects because the agent does not fully consider the potential for the shock to persist and because the policy occurs far in the future. Increasing k at first generates pessimism as the negative impact of the persistent shock dominates the positive impact of future policy. For sufficiently high k the policy effect ameliorates, and can dominate the pessimism. The net result is that some low level-k forecasts and some high level-k forecast generate very similar predictions, which removes the incentive under the replicator dynamics for low-level reasoners to revise upward.

Figure 7 illustrates this non-monotonicity. We assume $t = 0$, i.e. the stagnation state is initiated, and we use the Eggertsson et al. (2003) calibration of the model with $\xi = 0.99$, $\kappa = 0.2$, $\sigma = 0.5$, $\delta = 0.1$, and $r_N = 0.01$ with a large shock of $r_S = -0.01$. The level-0 forecast is set to $a^x = a^\pi = 0$. Each panel shows the same the progression of level-k inflation forecasts associated with three different forward guidance promises (i.e. values of ν) . The forward guidance promise ν^* corresponds to optimal promise under RE when ψ_x equal to the welfare theoretic value of 0.00254.[33] The solid lines illustrate non-monotonicity: inflation forecasts at first decrease in k and then increase (the same is true for the output gap forecast). Level-500 is approximately RE in each case.

The four panels in Figure 7 are distinguished by the distribution of level-k types used to compute the realized value of inflation in period $t = 1$, and as indicated with dashed lines. For each ν, the level-k forecast with lowest absolute error is indicated by the large dot. The top panels show the results for uniform distributions of level-k types between 0 and 499, and 0 and 3, respectively. The left bottom panel shows the results for the distribution of types observed in our pooled T3×A2 and A3 experiments for the first announcement treatment, where we observed proportions of level-0, 1, 2, 3, and REE forecasts of 25.4%, 56.8%, 4.2%, 1.7%, and 11.9%, respectively. The last panel shows an extreme case where a weight 1/2 is placed on level-0 and a weight of 1/2 is placed on level-500.

In all four cases the optimal level-k is smaller than the largest k in use, and for large forward guidance promises (smaller ν) there are double crossing of level-k

[33] The welfare theoretic value is κ/θ, where θ is the elasticity of demand with respect to price faced by the monopolistically competitive firms in the economy.

Figure 7: Forward guidance under unified dynamics simulation

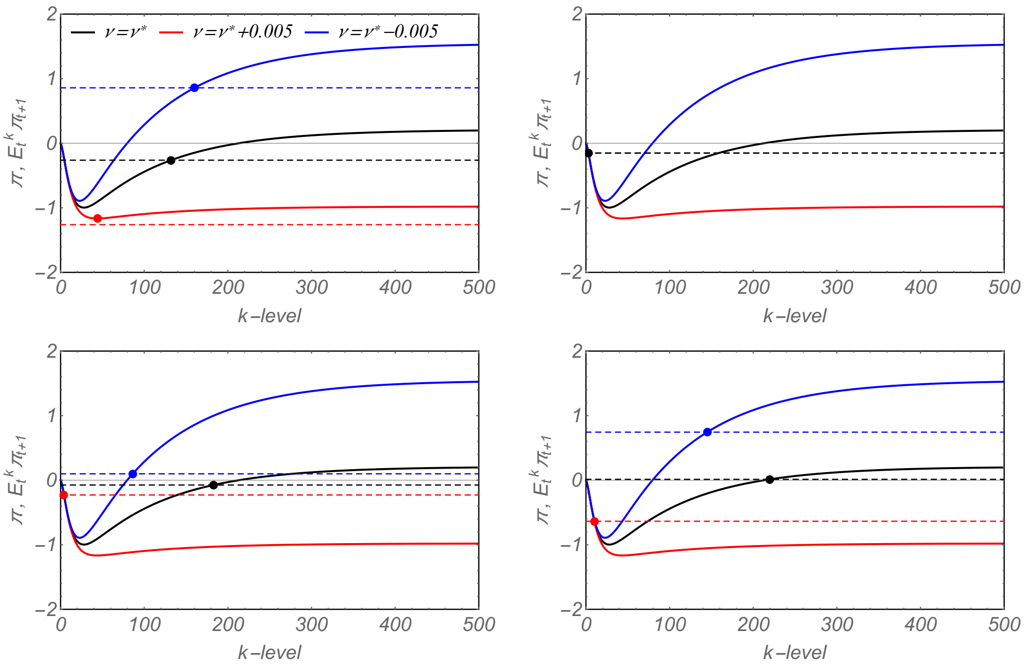

Notes: Each panel shows successive level-k deductions in the NK model at the ZLB in the shock state with differing forward guidance promises: ν. Dashed lines indicate the actual inflation outcome with level-0 beliefs at steady state and different proportions of level-k reasoners: uniform [0, 499], [0, 3], matched to round 20 from T3×A2 and A3 experiment, and level-0 and REE each 1/2. Comparisons between solid and dashed line provide the counterfactual that agents consider when revising strategies.

forecasts and realized inflation. This is significant because it reveals that low level-k and high level-k forecasts can both produce similarly small forecast errors: low level-k strategy do not lose many users even when the optimal level-k is high.

Figure 8 compares inflation and output dynamics of the model under RE (dashed black) and under unified dynamics (solid black), using the same calibration, including the initial distribution of level-k types, as in the bottom left panel of Figure 7. We set the learning gain to $\phi = 0.2$ and the replicator parameter to the relatively high value of $\alpha = 500$: this corresponds to a loss of nearly 90% of the users of a level-k strategy in a single period if the absolute forecast error for inflation is one percentage point. We note that by period 10, essentially no one is using the poorly-performing RE forecast; on the other hand, because low-level reasoning forecasts remain quite good, the collective use of the level-0, 1, 2, and 3 forecasts is still above 30%.

In the stagnation regime, unified dynamics leads to inflation that is below the RE value, and the discrepancy increases over time. Output is initially higher than under RE because of the myopia of low level-k reasoners. However, over time it too deteriorates through a combination of increasingly pessimistic level-k forecasts as k rises and because of revisions to the level-0 forecasts (red dotted line). During the forward guidance period itself (i.e. state F) inflation increases over time, but remains below RE, and the output gap follows a similar pattern.

Figure 8: Forward guidance under unified dynamics simulation

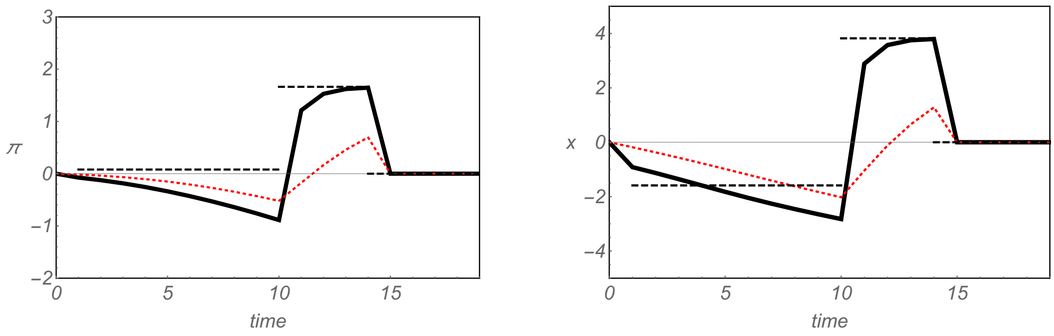

Notes: Unified dynamics (solid black) compared to RE (dashed black) for a realized shock of 10 quarters and a realized zero interest rate policy of four quarters. The red dashed line shows the path of level-0 beliefs. Parameters given the text.

7.4 Optimal forward guidance policy under the unified dynamic

Finally, we use simulations to solve for optimal forward guidance policy under the unified dynamics. Details are given in the online Appendix. We again use the initial distribution of types observed in our experiment in the first announcement round of the T3×A2 and A3 treatments, and we explore how optimal forward guidance policy changes, relative to RE, for different combinations of ϕ and α. Table 6 summarizes the results.

Optimal forward guidance policy under unified dynamics requires larger forward guidance promises for the broad ranges of α and ϕ we considered.[34] This is similar to the results found by Eusepi, Gibbs and Preston (2022). As α is reduced, the proportions of level-k reasoners are less responsive to forecast errors, which raises the importance of adaptive learning as the key constraint for policymakers. The power of forward guidance is diminished and the risk of the economy deteriorating in the event of a prolonged downturn is acute. Therefore larger forward guidance is indicated.

7.5 Discussion

Our model addresses observations made by Farhi and Werning (2019), who write,

> We close with the following general observations regarding level-k modeling. First, our analysis can either be interpreted as representing the impact

[34]When α is very large (two orders of magnitude larger than the largest value we considered for our simulations) smaller forward guidance promises can be optimal. We also observed smaller optimal forward guidance associated with $\phi = .2$ and $\alpha = 500$, though we suspect this is a result of numerical imprecision.

Table 6: Optimal forward guidance policy

		α			
ϕ		0	50	500	50000
0	ν^*	0.1730	0.1720	0.1690	0.1695
0.1	ν^*	0.1715	0.1710	0.1756	0.1760
0.2	ν^*	0.1710	0.1700	0.1765	0.1778
0.3	ν^*	0.1700	0.1695	0.1750	0.1780
REE	ν^*	0.1756	0.1756	0.1756	0.1756

Notes: Optimal forward guidance policies for different parameterizations of unified dynamics. Lower ν corresponds to a longer duration of promised interest rates. Remaining parameters given in the text.

> effect of interest rate changes or the dynamic effects in a world in which agents do not respond when they see realizations that differ from what they expect. Modeling how level-k agents react when they see unexpected realizations would require some hybrid of level-k reasoning and learning that is beyond the scope of the current paper, but is an interesting area for future research.

Our model is naturally viewed as a hybrid of level-k reasoning and learning. Our agents *do* respond when they see realizations different than they expect. Indeed they respond through two channels: by adjusting level-0 forecasts in light of new data, and by revising reasoning levels in light of performance. In the event of a stagnation shock, these channels induce an extended period of low-level reasoning, mitigating the impact of forward guidance.

8 Related Literature

Unified dynamics models how economic agents learn over time to forecast key endogenous aggregate variables. We have studied its implications for convergence to and departures from the rational expectations benchmark, and examined its implications for macroeconomic policy. In addition to the level-k literature, and the related cognitive hierarchy and reflective equilibrium approaches discussed in the Introduction, the unified model draws on several well-established literatures including,eductive stability, adaptive learning, and behavioral models.

The eductive approach analyzed in Guesnerie (1992) and Guesnerie (2002) examines the inherent difficulty for rational agents, who fully know the structural model, to coordinate on REE. These papers showed that coordination on an REE requires extremely strong common knowledge assumptions, not only of the structure but also of the rationality of other agents; even then coordination on an REE is only possible if the structure satisfies certain "eductive stability" conditions. These are closely related to the iterative expectational stability conditions developed in Evans (1985), a connection developed explicitly in Evans and Guesnerie (1993). In our unified model, level-k forecasts are obtained analogously, starting from level-0 forecasts, using iterations based on the structure.

Our level-0 learning assumption is based on the adaptive learning (AL) literature developed in Bray and Savin (1986), Marcet and Sargent (1989), Evans (1989) and Evans and Honkapohja (2001). AL is a versatile technique that has been applied in both nonexperimental and experimental settings. For a wide range of models it is known that agents acting as econometricians, and using least squares to update the coefficients of their forecast rules, can learn over time to have RE.[35] Importantly, under AL agents do not need to know or estimate structural parameters of the model, but simply regress the variable(s) they need to forecast on an intercept and relevant regressors, updating coefficients over time. AL thus provides a natural level-0 benchmark for deriving level-k forecasts.[36]

To complete our model we draw on the behavioral heterogeneous expectations literature Brock and Hommes (1997), De Grauwe (2012) and Hommes (2013), which, in a variety of macroeconomic settings, considers ex ante homogeneous agents selecting from a menu of forecast rules, resulting in ex post heterogeneity of forecasts. In our setting the menu includes the full set of level-k forecasts, and agents can raise or lower depth of reasoning based on recent forecast performance.

Our model shares elements with the Reflective Equilibrium notion proposed by García-Schmidt and Woodford (2019), which features heterogeneous level-k reasoners. Their approach develops a continuous version of level-k forecasts parameterized by a finite "degree of reflection" n, which corresponds to the mean level of reasoning. In their set-up RE would be obtained as $n \to \infty$, but they argue that a finite degree of reflection is more realistic, yielding a reflective equilibrium of degree n. As García-Schmidt and Woodford (2019) note, this approach is similar to the "calculation equilibrium" analyzed in Evans and Ramey (1992) and Evans and Ramey (1998), in which agents revise expectations of future paths recursively, with increased calculation costs from higher levels of recursion balanced against reduced forecast errors, leading to deviations of the calculation equilibrium path from REE. Presumably the García-Schmidt and Woodford (2019) finite degree

[35]Indeed, as stressed, for example, by Sargent (2008), REE are most plausibly viewed as an emergent outcome from learned behavior.

[36]Even generally well-behaved models can fail to be eductively stable. Evans, Guesnerie and McGough (2018) show that the standard Real Business Cycle model is not eductively stable even though it is stable under adaptive learning.

of reflection could be motivated by the additional mental effort or calculation costs associated with larger n. It would be straightforward for us to include such costs in the replicator component of our unified framework. We do not include these costs, mainly for simplicity, but also because we view that friction as less important than the core issue of the strategic uncertainty of decision-making.

Our framework also shares elements with models of rational inattention (e.g. Sims (2003) and Sims (2006)), and is supported by experimental evidence for sluggish discrete updating of beliefs as documented by Khaw, Stevens and Woodford (2017). We assume agents are inattentive with respect to their depth of reasoning when their forecasts are performing relatively well. This can lead to large forecast errors when the economy's structure changes – if these errors are large enough, subjects change their depth of reasoning.

Our LtFE shares important elements with the laboratory experiments of Fehr and Tyran (2008), Heemeijer, Hommes, Sonnemans and Tuinstra (2009), and Bao, Hommes, Sonnemans and Tuinstra (2012). Each study experimentally tests for convergence to an REE in an LtFE setting. Bao et al. (2012) study laboratory subjects' forecasts in settings with structural change similar to our announced structural change treatments. However in that paper subjects are not given the detailed structure of the model, and level-k forecasts are therefore not studied. Using a pricing game, Heemeijer et al. (2009) find that negative feedback engenders stabiity while positive feedback can lead to endogeneous fluctuations.[37] Fehr and Tyran (2008) study speed of convergence in a pricing game with different feedback treatments, which they refer to as strategic substitutability ($\beta < 0$ in our setup) and strategic complementarity ($0 < \beta < 1$). They argue that when $\beta < 0$, larger errors cause agents to update beliefs more quickly, leading to faster convergence. In contrast, under the unified model, additional forces associated with the distribution of level-k types and the magnitude and sign of β can slow or speed up convergence. In fact, if $\beta < -1$, large forecast errors may even prevent convergence rather than hasten its arrival.

Our study is also related to the experiments of Khaw, Stevens and Woodford (2019) and Anufriev, Duffy and Panchenko (2022), which both study forecasting tasks that nest a repeated beauty contest. Khaw et al. (2019) study forecasting with partial information and stochastic structural change following a Markov process, which is similar to our announced structural change treatments. Khaw et al. (2019) tests for level-k reasoning among participants and observe heterogeneous forecasts with different depths of reasoning, consistent with our findings.

In Anufriev et al. (2022), subjects must forecast two variables whose real-

[37]See also Sutan and Willinger (2009) for experimental work with negative feedback and level-k behavior.

izations are dependent on each other to capture more complicated expectational feedback environments. They compare their experimental data against a number of models that mix adaptive learning and level-k reasoning; both are necessary to fit the data. By contrast, our unified approach provides sharp predictions about revisions to depth of reasoning and the impact of anticipated events, and our experiment was designed to test these predictions.

9 Conclusion

The union of behavioral heterogeneity, adaptive learning, and level-k reasoning brings together three behavioral assumptions that enjoy wide experimental support. Level-k reasoning has been found to well describe how people form higher-order beliefs in wide variety of settings. We contribute to this literature by showing how level-k beliefs naturally fit with some of the most common forms of bounded rationality studied in macroeconomic environments. In addition, we provide a plausible way in which level-k beliefs may evolve over time in response to forecast errors and in response to adaptive learning through the level-0 forecast. A key finding is the persistence of low-level reasoners in environments with repeated structural change. This finding supports macroeconomic models that rely on low levels of reasoning to moderate general equilibrium effects.

Our experiment provides evidence for the key features of the unified model. We observe heterogeneous behavior consistent with level-k deductions as well as revisions to participants' depth of reasoning in line with the replicator dynamic. These results show how insights from beauty contest and cobweb model experiments extend to dynamic settings, and provide experimental support for the unified model to explain boundedly rational responses to announcements and hence to anticipated events.

The application of the unified theory to the forward guidance puzzle demonstrates the potentially wide applicability our approach to key issues in macroeconomics. The integration of adaptive learning, level-k reasoning and replicator dynamics explains the persistence of low reasoning levels following a stagnation shock, mitigating the otherwise unrealistic efficacy of monetary policy implemented via forward guidance.

The unified approach has a number of features that we find appealing. It naturally balances adaptive and eductive approaches to expectations formation by providing agents with some structural knowledge while also assuming they update beliefs over time as new data become available; it is amenable to theoretical analysis and yields both intuitive and surprising results; it is shown to be supported in experiments; and it is easily and naturally adapted to more general economic models.

References

Angeletos, George-Marios and Chen Lian, "Forward guidance without common knowledge," *American Economic Review*, 2018, *108* (9), 2477–2512.

— **and Karthik A Sastry**, "Managing Expectations: Instruments Versus Targets," *The Quarterly Journal of Economics*, 2021, *136* (4), 2467–2532.

Anufriev, Mikhail, John Duffy, and Valentyn Panchenko, "Learning in two-dimensional beauty contest games: Theory and experimental evidence," *Journal of Economic Theory*, 2022, *201*, 105417.

Arifovic, Jasmina, Stephanie Schmitt-Grohé, and Martin Uribe, "Learning to live in a liquidity trap," *Journal of Economic Dynamics and Control*, 2018, *89*, 120–136.

Bao, Te and John Duffy, "Adaptive versus eductive learning: Theory and evidence," *European Economic Review*, 2016, *83*, 64–89.

— **, Cars Hommes, Joep Sonnemans, and Jan Tuinstra**, "Individual expectations, limited rationality and aggregate outcomes," *Journal of Economic Dynamics and Control*, 2012, *36* (8), 1101–1120.

Bianchi, Francesco and Leonardo Melosi, "Dormant shocks and fiscal virtue," *NBER Macroeconomics Annual*, 2014, *28* (1), 1–46.

— **, Martin Lettau, and Sydney C Ludvigson**, "Monetary policy and asset valuation," *The Journal of Finance*, 2022, *77* (2), 967–1017.

Bilbiie, Florin O, "Optimal forward guidance," *American Economic Journal: Macroeconomics*, 2019, *11* (4), 310–345.

Bosch-Domenech, Antoni, Jose G Montalvo, Rosemarie Nagel, and Albert Satorra, "One, two,(three), infinity,...: Newspaper and lab beauty-contest experiments," *American Economic Review*, 2002, *92* (5), 1687–1701.

Branch, William A and Bruce McGough, "Replicator dynamics in a cobweb model with rationally heterogeneous expectations," *Journal of Economic Behavior & Organization*, 2008, *65* (2), 224–244.

Bray, Margaret M and Nathan E Savin, "Rational expectations equilibria, learning, and model specification," *Econometrica: Journal of the Econometric Society*, 1986, pp. 1129–1160.

Brock, William A and Cars H Hommes, "A rational route to randomness," *Econometrica: Journal of the Econometric Society*, 1997, pp. 1059–1095.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, "A cognitive hierarchy model of games," *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Costa-Gomes, Miguel A and Vincent P Crawford, "Cognition and behavior in two-person guessing games: An experimental study," *American Economic Review*, 2006, *96* (5), 1737–1768.

Duffy, John and Rosemarie Nagel, "On the robustness of behaviour in experimental 'beauty contest'games," *The Economic Journal*, 1997, *107* (445), 1684–1700.

Eggertsson, Gauti B et al., "Zero bound on interest rates and optimal monetary policy," *Brookings papers on economic activity*, 2003, *2003* (1), 139–233.

Eusepi, Stefano, Christopher Gibbs, and Bruce Preston, "Monetary policy trade-offs at the zero lower bound," *Available at SSRN 4069078*, 2022.

Evans, George W, "Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models," *The Quarterly Journal of Economics*, 1985, *100*, 1217–1233.

_ , "The Fragility of Sunspots and Bubbles," *Journal of Monetary Economics*, 1989, *23*, 297–317.

_ **and Bruce McGough**, "Adaptive Learning in Macroeconomics," *Oxford Research Encyclopedia of Economics and Finance*, 2020, *https://doi.org/10.1093/acrefore/9780190625979.013.508.*

_ **and** _ , "Agent-level Adaptive Learning," *Oxford Research Encyclopedia of Economics and Finance*, 2021, *https://doi.org/10.1093/acrefore/9780190625979.013.620.*

_ **and Garey Ramey**, "Expectation Calculation and Macroeconomic Dynamics," *American Economic Review*, 1992, *82*, 207–224.

_ **and** _ , "Calculation, Adaptation and Rational Expectations," *Macroeconomic Dynamics*, 1998, *2*, 156–182.

_ **and Roger Guesnerie**, "Rationalizability, Strong Rationality, and Expectational Stability," *Games and Economic Behaviour*, 1993, *5*, 632–646.

_ **and Seppo Honkapohja**, *Learning and expectations in macroeconomics*, Princeton University Press, 2001.

_ **, Roger Guesnerie, and Bruce McGough**, "Eductive Stability in Real Business Cycle Models," *The Economic Journal*, 05 2018, *129* (618), 821–852.

_ **, Seppo Honkapohja, and Kaushik Mitra**, "Anticipated fiscal policy and adaptive learning," *Journal of Monetary Economics*, 2009, *56* (7), 930–953.

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

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

Gaballo, Gaetano, "Eductive learning and the rationalizability of oligopoly games," *Economics Letters*, 2013, *120* (3), 401–404.

Gali, Jordi and Mark Gertler, "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, 1999, *44* (2), 195–222.

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

Gibbs, Christopher G, "Strongly Rational Routes to Randomness," *Working Paper, University of Sydney*, 2016.

_ **and Mariano Kulish**, "Disinflations in a model of imperfectly anchored expectations," *European Economic Review*, 2017, *100*, 157–174.

Goy, Gavin, Cars Hommes, and Kostas Mavromatis, "Forward guidance and the role of central bank credibility under heterogeneous beliefs," *Journal of Economic Behavior & Organization*, 2020.

Grauwe, Paul De, *Lectures on behavioral macroeconomics*, Princeton University Press, 2012.

Guesnerie, Roger, "An Exploration of the Eductive Justifications of the Rational-Expectations Hypothesis," *American Economic Review*, 1992, *82*, 1254–1278.

— , "Anchoring Economic Predictions in Common Knowledge," *Econometrica*, 2002, *70*, 439–480.

Heemeijer, Peter, Cars Hommes, Joep Sonnemans, and Jan Tuinstra, "Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation," *Journal of Economic Dynamics and Control*, 2009, *33* (5), 1052–1072. Complexity in Economics and Finance.

Ho, Teck-Hua, Colin Camerer, and Keith Weigelt, "Iterated dominance and iterated best response in experimental" p-beauty contests"," *The American Economic Review*, 1998, *88* (4), 947–969.

Hommes, Cars, "The heterogeneous expectations hypothesis: Some evidence from the lab," *Journal of Economic dynamics and control*, 2011, *35* (1), 1–24.

— , *Behavioral rationality and heterogeneous expectations in complex economic systems*, Cambridge University Press, 2013.

— , Joep Sonnemans, Jan Tuinstra, and Henk Van De Velden, "Learning in cobweb experiments," *Macroeconomic Dynamics*, 2007, *11* (S1), 8–33.

Jackson, Aaron L., "Disinflationary Boom Reversion," *Macroeconomic Dynamics*, September 2005, *9* (4), 489–515.

Jr., Robert E. Lucas, "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 1972, *4*, 103–124.

Kahneman, D., *Thinking fast and slow*, Farrar, Straus and Giroux, New York, 2011.

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

— , — , and — , "Adjustment dynamics during a strategic estimation task," 10 2019. "Working paper".

Kryvtsov, Oleksiy and Luba Petersen, "Central bank communication that works: Lessons from lab experiments," *Journal of Monetary Economics*, 2021, *117*, 760–780.

Marcet, Albert and Thomas J Sargent, "Convergence of least squares learning mechanisms in self-referential linear stochastic models," *Journal of Economic theory*, 1989, *48* (2), 337–368.

Mauersberger, Felix and Rosemarie Nagel, "Levels of reasoning in Keynesian Beauty Contests: a generative framework," in "Handbook of computational economics," Vol. 4, Elsevier, 2018, pp. 541–634.

Mokhtarzadeh, Fatemeh and Luba Petersen, "Coordinating expectations through central bank projections," *Experimental economics*, 2021, *24*, 883–918.

Muth, John F., "Rational Expectations and the Theory of Price Movements," *Econometrica*, 1961, *29*, 315–335.

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

— , Christoph Bühren, and Björn Frank, "Inspired and inspiring: Hervé Moulin and the discovery of the beauty contest game," *Mathematical Social Sciences*, 2017, *90*, 191–207. In honor of Hervé Moulin.

Preston, Bruce, "Learning About Monetary Policy Rules when Long-Horizon Expectations Matter," *International Journal of Central Banking*, 2005, *1* (2), 81–126.

Sargent, Thomas J., "Evolution and Intelligent Design," *American Economic Review*, 2008, *98*, 5–37.

Sethi, Rajiv and Reiner Franke, "Behavioural heterogeneity under evolutionary pressure: macroeconomic implications of costly optimisation," *The Economic Journal*, 1995, *105* (430), 583–600.

Sims, Christopher A, "Implications of rational inattention," *Journal of monetary Economics*, 2003, *50* (3), 665–690.

— , "Rational inattention: Beyond the linear-quadratic case," *American Economic Review*, 2006, *96* (2), 158–163.

Sutan, Angela and Marc Willinger, "Guessing with negative feedback: An experiment," *Journal of Economic Dynamics and Control*, 2009, *33* (5), 1123–1133. Complexity in Economics and Finance.

Weibull, Jörgen W, *Evolutionary game theory*, MIT press, 1997.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton, NJ: Princeton University Press, 2003.

For Online Publication

A1 Proofs

First, we formally establish our earlier contention that $\hat{k}$ is independent of level-0 beliefs and of the value of the constant γ.

Lemma A.1. *Fix β and weight system ω. Let $\hat{k}(\beta,\omega,a,\gamma)$ be the optimal sophistication level given the constant term γ and level-0 beliefs a. Then $\hat{k}(\beta,\omega,a,\gamma) = \hat{k}(\beta,\omega,1,0)$.*

Proof. Write $\mathcal{T}_\gamma(a,\omega,\beta)$ as the realized value of y given the datum (β,ω,a,γ). From equation (2) we have

$$\begin{aligned}\mathcal{T}_\gamma(a,\omega,\beta) - \frac{\gamma}{1-\beta} &= \gamma + \frac{\beta\gamma}{1-\beta}\sum_{k\geq 0}\omega_k - \frac{\beta\gamma}{1-\beta}\sum_{k\geq 0}\beta^k\omega_k + \beta a\sum_{k\geq 0}\beta^k\omega_k - \frac{\gamma}{1-\beta}\\ &= \frac{\gamma}{1-\beta} - \left(\frac{\gamma}{1-\beta}\right)\beta\sum_{k\geq 0}\beta^k\omega_k + \beta a\sum_{k\geq 0}\beta^k\omega_k - \frac{\gamma}{1-\beta}\\ &= \mathcal{T}_0\left(a - \frac{\gamma}{1-\beta},\omega,\beta\right).\end{aligned} \tag{A1}$$

Next, let $\phi(\beta,k,a,\gamma)$ be the forecast of a k-level agent. Then

$$\phi(\beta,k,a,\gamma) = \gamma\left(\frac{1-\beta^k}{1-\beta}\right) + \beta^k a.$$

Also, let $\phi^\varepsilon(\beta,k,a,\gamma) = |\phi(\beta,k,a,\gamma) - \mathcal{T}_\gamma(a,\omega,\beta)|$ be the associated forecast error.

Now observe that

$$\begin{aligned}\arg\min_{k\in\mathbb{N}}\phi^\varepsilon(\beta,k,a,0) &= \arg\min_{k\in\mathbb{N}}\left|a\beta^k - a\beta\sum_{n\geq 0}\beta^n\omega_n\right| = \arg\min_{k\in\mathbb{N}}|a|\left|\beta^k - \beta\sum_{n\geq 0}\beta^n\omega_n\right|\\ &= \arg\min_{k\in\mathbb{N}}\left|\beta^k - \beta\sum_{n>0}\beta^n\omega_n\right| = \arg\min_{k\in\mathbb{N}}\phi^\varepsilon(\beta,k,1,0).\end{aligned} \tag{A2}$$

Also, by (A1) we have that

$$\phi^\varepsilon(\beta,k,a,\gamma) = \phi^\varepsilon(\beta,k,a-\bar{y},0),$$

where $\bar{y} = \gamma(1-\beta)^{-1}$, so that

$$\arg\min_{k\in\mathbb{N}}\phi^\varepsilon(\beta,k,a,\gamma) = \arg\min_{k\in\mathbb{N}}\phi^\varepsilon(\beta,k,a-\bar{y},0). \tag{A3}$$

Putting (A2) and (A3) together yields

$$\arg\min_{k\in\mathbb{N}}\phi^\varepsilon(\beta,k,a,\gamma) = \arg\min_{k\in\mathbb{N}}\phi^\varepsilon(\beta,k,1,0),$$

which completes the proof. ■

Stability of unified dynamics. The strategy is to show that adaptive dynamics lead to convergence for any sequence of weights. Some notation is needed. Given a system of weights $\omega = \{\omega_i\}_{i\geq 0}$, let

$$\mathcal{T}_\gamma(a,\omega,\beta) = \gamma\left(1 + \frac{\beta}{1-\beta}\sum_{k\geq 0}(1-\beta^k)\omega_k\right) + \beta\sum_{k\geq 0}\beta^k\omega_k a \tag{A4}$$

Now fix any *sequence* of weight systems $\{\omega_t\}_{t\geq 0} = \left\{\{\omega_{it}\}_{i\geq 0}\right\}_{t\geq 0}$, and define the following recursion:

$$a_t = a_{t-1} + \phi\left(\mathcal{T}_\gamma\left(a_{t-1},\omega_{t-1},\beta\right) - a_{t-1}\right). \tag{A5}$$

We have the following result.

Lemma A.2. *Let $\phi \in (0,1]$.*

1. *If $|\beta| < 1$ then $a_t \to 0$.*
2. *If $\beta > 1$ then $|a_t| \to \infty$.*

Proof. First, observe that (A1) and (A5) imply

$$\begin{aligned} a_t - \frac{\gamma}{1-\beta} &= a_{t-1} - \frac{\gamma}{1-\beta} + \phi\left(\mathcal{T}_\gamma\left(a_{t-1},\omega_{t-1},\beta\right) - \frac{\gamma}{1-\beta} - \left(a_{t-1} - \frac{\gamma}{1-\beta}\right)\right) \\ &= a_{t-1} - \frac{\gamma}{1-\beta} + \phi\left(\mathcal{T}_0\left(a_{t-1} - \frac{\gamma}{1-\beta},\omega_{t-1},\beta\right) - \left(a_{t-1} - \frac{\gamma}{1-\beta}\right)\right), \end{aligned}$$

which shows that it suffices to prove the results for $\gamma = 0$. We drop the subscript on T.

Now assume $|\beta| < 1$, and observe that *for any* ω,

$$\left|\beta\sum_{k\geq 0}\beta^k\omega_k\right| \leq |\beta|\sum_{k\geq 0}|\beta^k|\omega_k \leq |\beta|\sum_{k\geq 0}|\beta|\,\omega_k \leq \beta^2. \tag{A6}$$

Next, write the recursion (A5) as

$$a_t = \left(1 - \phi\left(1 - \beta\sum_{k\geq 0}\beta^k\omega_{kt-1}\right)\right)a_{t-1} \equiv A_{t-1}a_{t-1}.$$

By equation (A6),

$$-1 < 1 - \phi(1+\beta^2) \leq A_{t-1} \leq 1 - \phi(1-\beta^2) < 1.$$

It follows that

$$|a_t| = \left(\prod_{n=1}^{t} A_{t-n}\right)|a_0| \to 0,$$

establishing item 1.

Now let $\beta > 1$. The same reasoning as in (A6), but with the inequalities reversed, yields

$$\beta \sum_{k \geq 0} \beta^k \omega_k \geq \beta^2.$$

It follows that

$$A_t \geq 1 - \phi + \phi\beta^2 = 1 + \phi(\beta^2 - 1) > 1,$$

and the result follows. ■

Proof of Theorem 1. The result is immediate: since Lemma A.2 holds for any sequence of weight systems, it holds in particular for whatever system of weights is produced by the unified dynamics. ■

Stability of the replicator dynamic. We begin with three lemmas.

Lemma A.3. *Suppose $\gamma = 0$.*

1. *If $|\beta| < 1$ then $k < \hat{k}(y)$ implies that there exists $\delta \in (0,1)$ such that $|y| < (1-\delta)|a\beta^k|$.*
2. *If $\beta > 1$ then $k < \hat{k}(y)$ implies that there exists $\delta > 0$ such $|y| > (1+\delta)|a\beta^k|$.*

Proof. Assume $|\beta| < 1$. If $|y| < |a\beta^{\hat{k}}|$ we are done, so assume $|a\beta^{\hat{k}}| \leq |y|$. Let $\delta = 1/2(1 - |\beta^{\hat{k}-k}|)$. We claim $2|y| < |a\beta^{\hat{k}}| + |a\beta^{\hat{k}-1}|$. Indeed, by the optimality of $\hat{k}$,

$$|y| - |a\beta^{\hat{k}}| = |y - a\beta^{\hat{k}}| < |y - a\beta^{\hat{k}-1}| = |a\beta^{\hat{k}-1}| - |y|.$$

Thus we compute

$$\begin{aligned} |y| &< \frac{1}{2}\left(|a\beta^{\hat{k}}| + |a\beta^{\hat{k}-1}|\right) \leq \frac{1}{2}\left(|a\beta^{\hat{k}}| + |a\beta^{k}|\right) \\ &= \frac{1}{2}\left(|\beta^{\hat{k}-k}| + 1\right)|a\beta^k| = (1-\delta)|a\beta^k|. \end{aligned}$$

Now assume $\beta > 1$. We may also assume, without loss of generality, that $a > 0$. Let $\delta = 1/2(|\beta^{\hat{k}-k}| - 1)$. If $y > a\beta^{\hat{k}}$ we are done, so assume $a\beta^{\hat{k}} \geq y$. It follows that

$$a\beta^{\hat{k}} \geq y > \frac{a}{2}\left(\beta^{\hat{k}} + \beta^k\right) = \frac{1}{2}\left(\beta^{\hat{k}-k} + 1\right)a\beta^k = (1+\delta)a\beta^k,$$

where the second inequality follows from the definition of $\hat{k}$. ■

Lemma A.4. *Let $\gamma = 0$ and $\{y_t\}_{t \geq 1}$ be generated by the replicator, initialized with weights $\{\omega_{n0}\}_{n \in \mathbb{N}}$ and beliefs a. Let $\check{k} \geq 1$ and suppose there exists $N > 0$ such that $t \geq N$ implies $\hat{k}(y_t) > \check{k}$. Then $\lim_{t \to \infty} \omega_{nt} = 0$ for all $n \leq \check{k}$.*

Proof. Let $t \geq N$. First suppose $|\beta| < 1$. Since $\hat{k}(y_t) > \check{k}$, it follows from Lemma A.3 that $(1-\delta)|a\beta^{\check{k}}| > |y_t|$, for some $\delta \in (0,1)$. Thus $n \leq \check{k}$ implies

$$|a\beta^n - y_t| \geq |a\beta^n| - |y_t| > |a\beta^n| - (1-\delta)|a\beta^{\check{k}}| > 0.$$

Using this estimate in the replicator yields, and that $r' > 0$, we have, for $s \geq 1$,

$$\begin{aligned}\omega_{nt+s} &= (1 - r(|a\beta^n - y_{t+s-1}|))\,\omega_{nt+s-1} \\ &< \left(1 - r\left(|a\beta^n| - (1-\delta)|a\beta^{\check{k}}|\right)\right)\omega_{nt+s-1} \\ &< \left(1 - r\left(|a\beta^n| - (1-\delta)|a\beta^{\check{k}}|\right)\right)^s \omega_{nt-1}.\end{aligned}$$

Because $r(0) \geq 0$ it follows that $\omega_{nt+s} \to 0$ as $s \to \infty$.

Now suppose $\beta > 1$, and assume, without loss of generality, that $a > 0$. Since $\hat{k}(y_t) > \check{k}$, it follows from Lemma A.3 that $a\beta^{\check{k}}(1+\delta) < y_t$. Thus $n \leq \check{k}$ implies

$$|a\beta^n - y_t| = y_t - a\beta^n \geq (1+\delta)a\beta^{\check{k}} - a\beta^n > 0.$$

The argument now proceeds analogously to the case $|\beta| < 1$. ■

Lemma A.5. *If x_n is an integer sequence and $\liminf x_n = x < \infty$ then there exists $N > 0$ such that $n \geq N$ implies $x_n \geq x$.*

Proof. The result is trivial if $x = -\infty$ so assume otherwise. Let $\hat{x}_k = \inf_{n \geq k} x_n$. Then $\hat{x}_k$ is a non-decreasing integer sequence converging to x. Now simply choose N so that $|\hat{x}_N - x| < 1$. ■

We are now ready to prove the main result.

Proof of Theorem 2. By Lemma A.1 we may assume $\gamma = 0$. To thin notation, let $\hat{k}_t = \hat{k}(y_t)$. It is helpful to introduce the relation $\succ$: for $y \in \mathbb{R}$ and $m(y), n(y) \in \mathbb{N}$, write $m(y) \succ n(y)$ when the level-m forecast is superior to the level-n forecast, i.e.,

$$m(y) \succ n(y) \iff |y - a\beta^{m(y)}| < |y - a\beta^{n(y)}|.$$

Now set $\tilde{k} = \liminf \hat{k}_t$.

We consider the cases $\beta > 1$ and $|\beta| < 1$ separately, however, we note that for each case it suffices to show $\tilde{k} = \infty$. To see this, first consider the case $\beta > 1$, and note that without loss of generality we may assume $a > 0$. Let $\Delta > 0$ and pick m so that $a\beta^m > \Delta$. Since $\tilde{k} = \infty$ it follows that $\hat{k}_t \to \infty$, so pick $\hat{t}$ so that $t \geq \hat{t} \implies \hat{k}_t > m$. Finally, for $n \geq 1$ let $\Omega_t^l(n) = \sum_{k<n} \omega_{kt}$, and note that, by Lemma A.4, $\tilde{k} = \infty$ implies $\Omega_t^l(n) \to 0$ as $t \to \infty$. Thus

$$\lim_{t\to\infty} y_t = \lim_{t\to\infty} a\beta \sum_{n\in\mathbb{N}} \beta^n \omega_{nt} \geq \lim_{t\to\infty} \left(1 - \Omega_t^l(m)\right) a\beta^{m+1} = a\beta^{m+1} > \Delta.$$

Now suppose $|\beta| < 1$. By Lemma A.4, if $\hat{k}_t \to \infty$ then all the weights are driven to zero. If all the weights are driven to zero then $y_t \to 0$: indeed, writing, $\omega_t^{\max} = \max_{i\in\mathbb{N}} \omega_{it}$, we have

$$|y_t| = \left| a\beta \sum_{n\in\mathbb{N}} \omega_{nt}\beta^n \right| \leq \omega_t^{\max} |a\beta| \sum_{n\in\mathbb{N}} |\beta^n| \to 0,$$

since $\omega_t^{\max} \to 0$ as $t \to \infty$.

Our proof strategy is to assume $\tilde{k} < \infty$ and derive a contradiction. To this end, it suffices to find some $M > 0$ so that $t \geq M$ implies the existence of $m(y_t) > \tilde{k}$ with $m(y_t) \succ \tilde{k}$, as this contradicts the definition of $\tilde{k}$ as the limit infimum of the $\hat{k}_t$.

First, the easy case: $\beta > 1$; and again assume $a > 0$. Then $y_t \geq (1 - \Omega_t^l(\tilde{k}))a\beta^{\tilde{k}+1}$, and, by Lemmas A.4 and A.5, $\lim_{t\to\infty} \Omega_t^l(\tilde{k})) = 0$. It follows that eventually, $\tilde{k}+1 \succ \tilde{k}$, which is the desired contradiction.

Now, assume $|\beta| < 1$, and let N be chosen as in Lemma A.5. The desired contradiction is developed in three steps.

Step 1. We establish the following claim:

Claim. *Given $\varepsilon > 0$ there exists $\mathcal{M}(\varepsilon) > 0$ so that $t \geq \mathcal{M}(\varepsilon) \geq N$ implies $|y_t| < |a\beta|^{\tilde{k}+1}(1+\varepsilon)$.*

Proof of claim. We know that for all $t \geq N$ we have $\hat{k}_t \geq \tilde{k}$. It follows that, for $\tilde{k} \geq 1$,

$$\begin{aligned} |y_t| &\leq |a\beta| \sum_{k<\tilde{k}} |\beta|^k \omega_{kt} + |a\beta| \sum_{k\geq\tilde{k}} |\beta|^k \omega_{kt} \\ &< |a\beta|\Omega_t^l\left(\tilde{k}\right) + |a||\beta|^{\tilde{k}+1}\left(1 - \Omega_t^l\left(\tilde{k}\right)\right). \end{aligned}$$

By Lemma A.4 we have that $\Omega_t^l\left(\tilde{k}\right) \to 0$ as $t \to \infty$, which establishes the claim.

Step 2. We now prove the result when $0 < \beta < 1$. Choose $2\varepsilon < \beta^{-1} - 1$ so that

$$(1+\varepsilon)\beta^{n+1} < \frac{1}{2}\left(\beta^{n+1} + \beta^n\right).$$

Let $\mathcal{M}(\varepsilon) = \mathcal{M}$ be chosen as in Step 1, and assume $t \geq \mathcal{M}$. There are two cases.

Case 1: $a > 0$. It follows that $y_t > 0$. Then

$$0 < y_t < a\beta^{\tilde{k}+1}(1+\varepsilon) < \frac{1}{2}\left(a\beta^{\tilde{k}+1} + a\beta^{\tilde{k}}\right),$$

which implies that $\tilde{k}+1 \succ \tilde{k}$, the desired contradiction.

Case 2: $a < 0$. In this case $y_t < 0$. Then

$$0 > y_t > a\beta^{\tilde{k}+1}(1+\varepsilon) > \frac{1}{2}\left(a\beta^{\tilde{k}+1} + a\beta^{\tilde{k}}\right),$$

which implies that $\tilde{k}+1 \succ \tilde{k}$, the desired contradiction.

Step 3. Finally, we prove the result when $-1 < \beta < 0$. Choose $\varepsilon < (2|\beta|)^{-1}(1 - |\beta|)^2$ and choose $\mathcal{M}(\varepsilon)$ as in Step 1. Now notice that

$$1 + \varepsilon < (2|\beta|^{(n+1)})^{-1}\left(|\beta|^n + |\beta|^{n+2}\right),$$

for any $n \geq 1$. It follows that

$$\begin{aligned} 2|\beta|^{\tilde{k}+1}(1+\varepsilon) &< |\beta|^{\tilde{k}+2} + |\beta|^{\tilde{k}}, \text{ or} \\ 0 < |\beta|^{\tilde{k}+1}(1+\varepsilon) - |\beta|^{\tilde{k}+2} &< |\beta|^{\tilde{k}} - |\beta|^{\tilde{k}+1}(1+\varepsilon). \end{aligned} \tag{A7}$$

Let $t \geq M$. There are two cases.

Case 1: $\hat{k}_t \neq \tilde{k} \pmod 2$. In this case $\text{sign}(y_t) = -\text{sign}(a\beta^{\tilde{k}})$, whence $\tilde{k}+1 \succ \tilde{k}$.

Case 2: $\hat{k}_t = \tilde{k} \pmod 2$. If $y_t < 0$ then $a\beta^{\tilde{k}}$ is negative. Next, note that if $y_t \geq a\beta^{\tilde{k}+2}$ then $\tilde{k}+2 \succ \tilde{k}$, which is a contradiction. Thus

$$a\beta^{\tilde{k}} < -|a\beta^{\tilde{k}+1}|(1+\varepsilon) < y_t < a\beta^{\tilde{k}+2} < 0,$$

where the first inequality follows from (A7). Thus

$$\begin{aligned}
|a\beta^{\tilde{k}+2} - y_t| &< \left| a\beta^{\tilde{k}+2} + |a\beta^{\tilde{k}+1}|(1+\varepsilon) \right| \\
&= |a\beta^{\tilde{k}+1}|(1+\varepsilon) - |a\beta^{\tilde{k}+2}| \\
&= |a| \left(|\beta^{\tilde{k}+1}|(1+\varepsilon) - |\beta^{\tilde{k}+2}| \right) \\
&< |a| \left(|\beta^{\tilde{k}}| - |\beta^{\tilde{k}+1}(1+\varepsilon)| \right) \\
&= |a\beta^{\tilde{k}}| - |a\beta^{\tilde{k}+1}(1+\varepsilon)| \\
&< |a\beta^{\tilde{k}}| - |y_t| = \left| a\beta^{\tilde{k}} - y_t \right|,
\end{aligned}$$

which implies $\tilde{k}+2 \succ \tilde{k}$.

Now suppose $y_t > 0$, so that $a\beta^{\tilde{k}}$ is positive. Thus

$$a\beta^{\tilde{k}} > |a\beta^{\tilde{k}+1}|(1+\varepsilon) > y_t > a\beta^{\tilde{k}+2} > 0,$$

where the reasoning is as above. Thus

$$\begin{aligned}
|y_t - a\beta^{\tilde{k}+2}| &< \left| |a\beta^{\tilde{k}+1}|(1+\varepsilon) - a\beta^{\tilde{k}+2} \right| \\
&= |a\beta^{\tilde{k}+1}|(1+\varepsilon) - |a\beta^{\tilde{k}+2}| \\
&= |a| \left(|\beta^{\tilde{k}+1}|(1+\varepsilon) - |\beta^{\tilde{k}+2}| \right) \\
&< |a| \left(|\beta^{\tilde{k}}| - |\beta^{\tilde{k}+1}(1+\varepsilon)| \right) \\
&= |a\beta^{\tilde{k}}| - |a\beta^{\tilde{k}+1}(1+\varepsilon)| \\
&< |a\beta^{\tilde{k}}| - |y_t| = \left| a\beta^{\tilde{k}} - y_t \right|,
\end{aligned}$$

so that $\tilde{k}+2 \succ \tilde{k}$, completing the proof of step 3. ■

Proof of Theorem 3. Lemma A.2 establishes items 1 and 2, and so we focus here only on item 3. Also, as demonstrated in the proof of Lemma A.2, we may assume $\gamma = 0$. We recall the notation $\Omega = \dot{\cup}_n \Delta^n$ and $\psi_\beta : \Omega \to \mathbb{R}$, given by $\psi_\beta(\omega) = \beta \sum_k \beta^k \omega_k$, and that Ω is endowed with the direct-limit topology.

The dynamic system for a_t may be written

$$a_t = (1 - \phi + \phi\psi_\beta(\omega))\, a_{t-1} \equiv A(\beta, \omega, \phi) a_{t-1}.$$

It follows that $|A(\beta,\omega,\phi)| < 1 \implies a_t \to 0$ and $|A(\beta,\omega,\phi)| > 1 \implies |a_t| \to \infty$.

We compute

$$\begin{aligned}|A(\beta,\omega,\phi)| < 1 &\iff -1 < 1-\phi+\phi\psi_\beta(\omega) < 1 \iff 1-2\phi^{-1} < \psi_\beta(\omega) < 1, \text{ and}\\ |A(\beta,\omega,\phi)| > 1 &\iff 1-\phi+\phi\psi_\beta(\omega) < -1 \text{ or } 1-\phi+\phi\psi_\beta(\omega) > 1\\ &\iff \psi_\beta(\omega) < 1-2\phi^{-1} \text{ or } \psi_\beta(\omega) > 1.\end{aligned}$$

This completes the proof of items 3(a) - 3(c).

To establish item 3(d) we start by showing that ψ_β is continuous. Let ψ_β^n be the restriction of ψ_β to $\Delta^n \subset \Omega$. It suffices to show that $\psi_\beta^n : \Delta^n \to \mathbb{R}$ is continuous for each $n \in \mathbb{N}$. To see this, let $U \subset \mathbb{R}$ be open. Then

$$\psi_\beta^{-1}(U) = \cup_n \left(\psi_\beta^{-1}(U) \cap \Delta^n\right) = \cup_n \left(\left(\psi_\beta^n\right)^{-1}(U) \cap \Delta^n\right) = \cup_n \left(\left(\psi_\beta^n\right)^{-1}(U)\right).$$

Assuming $\psi_\beta^n : \Delta^n \to \mathbb{R}$ is continuous, we have that $\left(\psi_\beta^n\right)^{-1}(U)$ is open in Δ^n, whence open in Ω. Thus $\psi_\beta^{-1}(U)$ is a union of open sets in Ω, which establishes the continuity of ψ_β.

Next we demonstrate surjectivity of ψ_β. Let $z \in \mathbb{R}$. Since $\beta < -1$ we can find an $n \in \mathbb{N}$ with $n \geq 1$ so that $\beta^{2n+1} < z < \beta^{2n}$. By continuity there is $\varepsilon \in (0, 1/2)$ such that

$$(1-\varepsilon)\beta^{2n+1} + \varepsilon\beta^{2n} < z < \varepsilon\beta^{2n+1} + (1-\varepsilon)\beta^{2n}.$$

For $\alpha \in (0,1)$ let $\omega(\alpha) \in \Delta^{2n+1} \subset \Omega$ be given by

$$\omega_k(\alpha) = \begin{cases} \alpha & \text{if } k = 2n+1 \\ 1-\alpha & \text{if } k = 2n \\ 0 & \text{else} \end{cases}$$

and note that $\alpha \to \omega_\alpha$ continuously maps $(0,1)$ into Δ^{2n+1}, whence into Ω. Let $\Psi_\beta : (0,1) \to \mathbb{R}$ be $\Psi_\beta(\alpha) = \psi_\beta(\omega(\alpha))$. It follows that Ψ_β is continuous and

$$\Psi_\beta(\varepsilon) = (1-\varepsilon)\beta^{2n+1} + \varepsilon\beta^{2n} < z < \varepsilon\beta^{2n+1} + (1-\varepsilon)\beta^{2n} = \Psi_\beta(1-\varepsilon)$$

By the intermediate value theorem there is an $\alpha \in (\varepsilon, 1-\varepsilon)$ so that $z = \Psi_\beta(\alpha) = \psi_\beta(\omega(\alpha))$, which establishes surjectivity.

Now let

$$\begin{aligned}\Omega_s &= \psi_\beta^{-1}\left((1-2\phi^{-1}, 1)\right)\\ \Omega_u &= \psi_\beta^{-1}\left((-\infty, 1-2\phi^{-1}) \cup (1,\infty)\right).\end{aligned}$$

Both sets are open by the continuity of ψ_β, and from items 3(a) and 3(b) we have that $\omega \in \Omega_s$ implies $y_t \to \bar{y}$ and $\omega \in \Omega_u$ implies $|y_t| \to \infty$. Thus parts (i) and (ii) of item 3(d) are established.

Finally, let $\Omega_0 = \Omega \setminus (\Omega_s \cup \Omega_u)$. We must show that Ω_0 is no-where dense, i.e. that the interior of the closure of Ω_0 is empty. To this end, notice that

$$\Omega_0 = \psi_\beta^{-1}(\{-1\}) \,\dot\cup\, \psi_\beta^{-1}(\{1\}) \equiv \Omega_0^- \,\dot\cup\, \Omega_0^+.$$

Since ψ_β is continuous, it follows that $\Omega_0^\pm$ are closed. Since no-where denseness is closed under finite unions, it suffices to show that the interiors of $\Omega_0^\pm$ are empty. Thus let $\omega \in \Omega_0^+$. Let $N \in \mathbb{N}$ so that $\omega \in \Delta^N$. Since $\beta < -1$ and $\psi_\beta(\omega) = 1$ there is an even $n \in \mathbb{N}$ and an odd $m \in \mathbb{N}$, with $n, m \leq N$ and such that $\omega_n, \omega_m \neq 0$.

For $k \in \mathbb{N}$ with $k \geq 2$, define $\omega^k \in \Delta^N \subset \Omega$ as follows:

$$\omega_i^k = \begin{cases} (1-k^{-1})\omega_n & \text{if } i = n \\ \omega_m + k^{-1}\omega_n & \text{if } i = m \\ \omega_i & \text{else} \end{cases}$$

Note that ω^k is the same weight system as ω except that some of the weight associated with the positive forecast β^n is shifted to the negative forecast β^m. Because the model itself has negative feedback, this means that the implied value of y is larger for weight system ω^k than it is for weight system ω. More formally, $k \geq 2$ implies that $\psi_\beta(\omega^k) > 1$, which implies that $\omega^k \in \Omega_u$. Now notice that, as a sequence in Δ^N, we have $\omega^k \to \omega$. Owing to the construction of the direct-limit topology, we have that $\omega^k \to \omega$ in Ω as well. Thus, given an arbitrary element $\omega \in \Omega_0^+$ we have constructed a sequence in Ω_u converging to it, and since $\Omega_u \cap \Omega_0^+$ is empty, we conclude that ω is not in the interior of Ω_0^+. So the interior of Ω_0^+ is empty, and since Ω_0^+ is closed, we conclude that Ω_0^+ is nowhere dense. The same argument applies to Ω_0^-, which shows that $\Omega_0 = \Omega_0^- \dot\cup \Omega_0^+$ is no-where dense. ■

Full statement and proof of Proposition 1. Recall from (4) that $\hat{k}$ is defined explicitly as a function of y_t. However, both y_t and $E_{t-1}^k y_t$ are affine functions of level-0 beliefs a. In particular, if $\gamma = 0$ then

$$\hat{k}(a) = \min \arg\min_{k \in \mathbb{N}} |\beta^k a - \beta \textstyle\sum_k \omega_k a|, \tag{A8}$$

which further implies that $\hat{k}$ is independent of a. It is straightforward to show this result continues to hold with $\gamma \neq 0$, and, in fact, $\hat{k}$ is independent of the value of γ. Thus, we may view $\hat{k} = \hat{k}(\beta, \omega)$. We have the following result.

Proposition 1′ (Optimal forecast levels)**.** *Let $K \geq 1$ and $\omega^K = \{\omega_n\}_{n=0}^K$ be a weight system with weights given as $\omega_n = (K+1)^{-1}$. Let $\hat{k} = \hat{k}\left(\beta, \omega^K\right)$.*

1. *Suppose $0 < \beta < 1$.*
   (a) $K \to \infty \implies \hat{k} \to \infty$ *and* $\hat{k}/K \to 0$.
   (b) $\beta \to 1^- \implies \hat{k} \to \begin{cases} \frac{K}{2} + 1 & \text{if } K \text{ is even} \\ \frac{K+1}{2} & \text{if } K \text{ is odd} \end{cases}$
   (c) $\beta \to 0^+ \implies \hat{k} \to \begin{cases} 1 & \text{if } K = 1 \\ 2 & \text{if } K \geq 2 \end{cases}$
2. *Suppose $-1 < \beta < 0$.*
   (a) $K \to \infty \implies \hat{k} \to \infty$ *and* $\hat{k}/K \to 0$.
   (b) $\beta \to 0^- \implies \hat{k} \to \begin{cases} 1 & \text{if } K = 1 \\ 3 & \text{if } K \geq 2 \end{cases}$
   (c) $\beta \to -1^+ \implies \hat{k} \to \infty$.
3. *Suppose $\beta < -1$*
   (a) $K \to \infty \implies \hat{k} \to \infty$ *and* $\hat{k}/K \to 1$

(b) $\beta \to -1^- \implies \hat{k} \to \begin{cases} 1 \text{ if } K \text{ is even} \\ 0 \text{ if } K \text{ is odd} \end{cases}$

(c) $\beta \to -\infty \implies \hat{k} \to K+1$.

Before proceeding to the proof, some preliminary work is required. By Lemma A.1 we may assume $\gamma = 0$ and $a = 1$. Recall from Section 4.2 our notation for uniform weights: for $K \in \mathbb{N}$, $\omega^K = \{\omega_n\}_{n=0}^K$ with $\omega_n = (K+1)^{-1}$. It follows that

$$y = \beta \textstyle\sum_k \beta^k \omega_k = \frac{\beta}{K+1} \sum_k \beta^k = \frac{\beta\left(1-\beta^{K+1}\right)}{(K+1)(1-\beta)} \equiv \psi\left(K,\beta\right).$$

When it does not impede clarity, we make the identifications $\hat{k} = \hat{k}\left(\beta, \omega^K\right)$ and $\psi = \psi\left(K,\beta\right)$.

It is helpful to define k^* as the *continuous counterpart* to $\hat{k}$. For $\beta > 0$ our definition for k^* corresponds to the first order condition for minimizing $(\beta^k - \psi(K,\beta))^2$ for $k \in \mathbb{R}_+$. However, care must be taken to accommodate $\beta < 0$. We define k^* as follows:

$$k^*\left(K,\beta\right) = \frac{\log\left(\psi\left(K,\beta\right)^2\right)}{\log(\beta^2)}. \tag{A9}$$

Of course if β, and hence ψ, are positive then we can dispense with the squared terms in the definition.

Now define $\lfloor\cdot\rfloor$ to be the usual floor function, i.e. for $x \in \mathbb{R}$, $\lfloor x \rfloor$ is the largest integer less than or equal to x. Define $\lfloor\cdot\rfloor_{\text{odd}}$ and $\lfloor\cdot\rfloor_{\text{even}}$ and the odd and even floors, respective, which take the obvious meaning, e.g. $\lfloor x \rfloor_{\text{even}}$ is the largest even integer less than or equal to x. Finally, $\lceil\cdot\rceil$, $\lceil\cdot\rceil_{\text{even}}$, and $\lceil\cdot\rceil_{\text{even}}$ have the analogous definitions. Define

$$k^*_{\text{low}} = \begin{cases} \lfloor k^* \rfloor & \text{if } 0 < \beta < 1 \\ \lfloor k^* \rfloor_{\text{odd}} & \text{if } -1 < \beta < 0 \text{ or if } \beta < -1 \textit{ and } \psi < \frac{1+\beta}{2} \\ \lfloor k^* \rfloor_{\text{even}} & \text{if } \beta < -1 \textit{ and } \psi > 0 \end{cases}$$

and define k^*_{high} analogously using the ceiling functions. The following result links k^* and $\hat{k}$.

Lemma A.6. *If $k^* \geq 0$ then $\hat{k} \in \{k^*_{low}, k^*_{high}\}$.*

Proof. We begin with the following observations on the parity of $\hat{k}$.[38] Recall that 0 is taken as even.

1. If $-1 < \beta < 0$ then $\hat{k}$ is odd.
2. If $\beta < -1$ and $\psi < \frac{1+\beta}{2}$ then $\hat{k}$ is odd.
3. If $\beta < -1$ and $\psi > 0$ then $\hat{k}$ is even.

These items may be established as follows. Note that $-1 < \beta < 0$ implies $\psi < 0$, whence there is an odd $n \in \mathbb{N}$ so that $\psi < \beta^n < 0$, making n superior to any even forecast level. If $\beta < -1$ and $\psi < \frac{1+\beta}{2}$ then the level 1 forecast is superior to any even forecast level. If $\beta < -1$ and $\psi > 0$ then the level 0 forecast is superior to any odd forecast level.

[38] The *parity* of $n \in \mathbb{N}$ is its equivalence class mod 2. Thus n and m have the same parity if they are either both even or both odd.

Next, note that $k^* < 0$ if and only if $-1 < \psi < 1$ and $\beta < -1$. Now, for $\alpha \in \mathbb{R}_+$ define $\phi(\alpha,\beta)$ as follows:

$$\phi(\alpha,\beta) = \begin{cases} (\beta^2)^{\frac{\alpha}{2}} & \text{if } \psi > 0 \\ \beta\,(\beta^2)^{\frac{\alpha-1}{2}} & \text{if } \psi < 0 \end{cases}$$

This function has the following properties:
- (a) If $\hat{k} \geq 1$ and if non-zero $k \in \mathbb{N}$ has the same parity as $\hat{k}$ then $\beta^k = \phi(k,\beta)$: in this way ϕ extends our notion of forecast level to all positive reals.
- (b) $\phi(k^*,\beta) = \psi$.

To establish item (a), first suppose $\hat{k}$ is even. Since $\hat{k} \geq 1$ it follows that $\psi > 0$. Let $k = 2m$ for $m > 0$. Then $\phi(k,\beta) = (\beta^2)^m = \beta^k$. Next suppose $\hat{k}$ is odd. Let $k = 2m+1$. If $0 < \beta < 1$ then $\psi > 0$, so that $\phi(k,\beta) = (\beta^2)^{\frac{2m+1}{2}} = \beta^{2m+1}$. Let $\beta < 0$. If $-1 < \beta < 0$ then $\psi < 0$. If $\beta < -1$ then $\hat{k}$ odd implies $\psi < 0$. Thus $k = 2m+1$ implies $\phi(k,\beta) = \beta(\beta^2)^m = \beta^{2m+1}$. To establish item (b), observe that $\psi > 0$ implies

$$\log\phi(k^*,\beta) = ({k^*}/{2})\log\beta^2 = ({1}/{2})\log\psi^2 = \log\psi$$

and $\psi < 0$ implies $\phi(k^*,\beta) < 0$, and

$$\log\left(-\phi(k^*,\beta)\right) = \log\left(\beta^2\right)^{\frac{1}{2}}\left(\beta^2\right)^{\frac{k^*-1}{2}} = \log\left(\beta^2\right)^{\frac{k^*}{2}} = ({k^*}/{2})\log\beta^2 = \log\left(-\psi\right).$$

We turn now to the body of the proof of Lemma A.6, in which we use the following notation: $k_1 \prec k_2$ if β^{k_1} is strictly inferior to β^{k_2} as a forecast of ψ. The strategy is as follows: show that $k < \lfloor k^* \rfloor \implies k \prec \lfloor k^* \rfloor$, and that $k > \lceil k^* \rceil$ implies that $k \prec \lceil k^* \rceil$, with floor and ceiling functions adjusted for parity as needed.

Case 1: $0 < \beta < 1$. Since $\psi < \beta$ in this case, we have that $k^* \geq 1$ and $\hat{k} \geq 1$. Also $\alpha > 0$ implies $\phi_\alpha(\alpha,\beta) < 0$. Thus if $k_1 < \lfloor k^* \rfloor$ and $k_2 > \lceil k^* \rceil$ then

$$\phi(k_1,\beta) > \phi(\lfloor k^* \rfloor,\beta) \geq \underbrace{\phi(k^*,\beta)}_{\psi} \geq \phi(\lceil k^* \rceil,\beta) > \phi(k_2,\beta).$$

Thus $k_1 \prec \lfloor k^* \rfloor$ and $k_2 \prec \lceil k^* \rceil$.

Case 2: $-1 < \beta < 0$. Since $\beta < \psi < 0$ in this case, we have that $k^* \geq 1$. Also $\alpha > 0$ implies $\phi_\alpha(\alpha,\beta) > 0$. Also $\psi < 0$ so that $\hat{k}$ is necessarily odd. Thus if $\lfloor k^* \rfloor_{\text{odd}} \geq 1$ and if k_i are odd with $k_1 < \lfloor k^* \rfloor_{\text{odd}}$ and $k_2 > \lceil k^* \rceil_{\text{odd}}$, then

$$\phi(k_1,\beta) < \phi(\lfloor k^* \rfloor_{\text{odd}},\beta) \leq \underbrace{\phi(k^*,\beta)}_{\psi} \leq \phi(\lceil k^* \rceil_{\text{odd}},\beta) < \phi(k_2,\beta).$$

Thus $k_1 \prec \lfloor k^* \rfloor_{\text{odd}}$ and $k_2 \prec \lceil k^* \rceil_{\text{odd}}$.

Case 3: $\beta < -1$ and $\psi < \frac{1+\beta}{2}$. Then $k^* \geq 1$ and $\hat{k}$ is odd. Also $\alpha > 1$ implies $\phi_\alpha(\alpha,\beta) < 0$. Thus if $\lfloor k^* \rfloor_{\text{odd}} > 1$ and if k_i are odd with $k_1 < \lfloor k^* \rfloor_{\text{odd}}$ and $k_2 > \lceil k^* \rceil_{\text{odd}}$, then

$$\phi(k_1,\beta) > \phi(\lfloor k^* \rfloor_{\text{odd}},\beta) \geq \underbrace{\phi(k^*,\beta)}_{\psi} \geq \phi(\lceil k^* \rceil_{\text{odd}},\beta) > \phi(k_2,\beta).$$

Thus $k_1 \prec \lfloor k^* \rfloor_{\text{odd}}$ and $k_2 \prec \lceil k^* \rceil_{\text{odd}}$.

Case 4: $\beta < -1$ and $\psi > 0$. Then $k^* \geq 0$ (by assumption) and $\hat{k}$ is even.

Also $\alpha > 0$ implies $\phi_\alpha(\alpha,\beta) > 0$. Thus if $\lfloor k^* \rfloor_{\text{even}} > 2$ and if k_i are even with $k_1 < \lfloor k^* \rfloor_{\text{even}}$ and $k_2 > \lceil k^* \rceil_{\text{even}}$, then

$$\phi(k_1,\beta) < \phi(\lfloor k^* \rfloor,\beta) \leq \underbrace{\phi(k^*,\beta)}_{\psi} \leq \phi(\lceil k^* \rceil_{\text{even}},\beta) < \phi(k_2,\beta).$$

Thus $\lfloor k^* \rfloor_{\text{even}} > 2$ implies $k_1 \prec \lfloor k^* \rfloor_{\text{even}}$ and $k_2 \prec \lceil k^* \rceil_{\text{even}}$. If $\lfloor k^* \rfloor_{\text{even}} = 2$ then

$$1 \equiv \beta^0 < \beta^2 = \phi(\lfloor k^* \rfloor_{\text{even}},\beta) \leq \underbrace{\phi(k^*,\beta)}_{\psi} \leq \phi(\lceil k^* \rceil_{\text{even}},\beta) < \phi(k_2,\beta).$$

If $\lfloor k^* \rfloor_{\text{even}} = 0 < k^*$ then

$$1 \equiv \beta^0 < \underbrace{\phi(k^*,\beta)}_{\psi} \leq \phi(\lceil k^* \rceil_{\text{even}},\beta) < \phi(k_2,\beta).$$

Finally, if $k^* = 0$ then $\hat{k} = k^*$. ■

We now turn to the proof of Proposition 1. We note that if $K = 0$ then $k^* = \hat{k} = 1$ regardless of the value of β, so this case is excluded.

Proof of Proposition 1. The arguments for the limits involving $K \to \infty$ will rely directly on the behavior of k^*. The arguments involving limits in β require additional analysis. Define

$$\Delta(k_1,k_2,\beta) = \left(\beta^{k_1} - \psi(\beta)\right)^2 - \left(\beta^{k_2} - \psi(\beta)\right)^2,$$

and note that $k_1 \prec k_2$ when $\Delta(k_1,k_2,\beta) > 0$ and $k_2 \prec k_1$ when $\Delta(k_1,k_2,\beta) < 0$, where the ordering here is as defined in the proof of Lemma A.6. The proof strategy for limiting values of β has three steps:

1. Compute the relevant limiting value of k^*.
2. Use Lemma A.6 to determine a finite set $\hat{\mathcal{K}}$ of possible limiting values for $\hat{k}$.
3. Expand Δ around the limiting value of β and use the expansion to pairwise compare the elements of the $\hat{\mathcal{K}}$.

A final comment before proceeding: Many of the arguments below include tedious symbolic manipulation, and we have relegated much of this work to Mathematica. Whenever Mathematica is relied upon to reach a conclusion, we state this reliance explicitly. As an example, the code used for the first result is included below. All code is available upon request.

Case 1: $0 < \beta < 1$. The following Mathematica code establishes that $K \to \infty$ implies $k^* \to \infty$ and $k^*/K \to 0$.

```
psi[K_, beta_] := beta/(K + 1) Sum[beta^(k - 1), {k, 1, K + 1}];
kstar[K_, beta_] := Log[psi[K, beta]^2]/Log[beta^2];
Module[{limK, limKk, assume},
  assume = {0 < beta < 1};
  limK = Limit[kstar[K, beta], K -> \[Infinity], Assumptions -> And @@ assume];
  limKk = Limit[kstar[K, beta]/K, K -> \[Infinity], Assumptions -> And @@ assume];
  Print["Limit of kstar as K -> infinity is " <> ToString@limK];
  Print["Limit of kstar/K as K -> infinity is " <> ToString@limKk];
  ];
```

Lemma A.6 then implies the same limits for $\hat{k}$, thus proving item 1(a).

Turning to item 1(b), using Mathematica, we find that $\beta \to 1^-$ implies $k^* \to K/2 + 1$. Suppose K is odd. It follows that β near (and below) 1 implies $\lfloor k^* \rfloor <$

$k^* < \lceil k^* \rceil$, whence

$$\hat{k} \in \{\lfloor k^* \rfloor, \lceil k^* \rceil\} = \left\{\frac{K+1}{2}, \frac{K+3}{2}\right\}.$$

Using Mathematica, we find that near $\beta = 1$,

$$\Delta\left(\frac{K+1}{2}, \frac{K+3}{2}, \beta\right) = \frac{1}{12}(K-1)(K+3)(\beta-1)^3 + \mathcal{O}\left(|\beta-1|^4\right),$$

so that when $K \geq 3$ and β is near and below 1, we conclude that $\Delta < 0$, so that $\hat{k} = {}^1\!/_2(K+1)$. When $K = 1$ a direct computation shows $\Delta = 0$, so that both $\lfloor k^* \rfloor$ and $\lceil k^* \rceil$ yield the same forecast. Our tiebreaker, then, chooses $\hat{k} = 1$.

Now suppose K is even. Then for β near and below 1 we know that k^* is near ${}^K\!/_2 + 1 \in \mathbb{N}$. Unfortunately, we do not know if k^* approaches its limit monotonically. Thus we can only conclude that for β near and below 1 we have

$$\hat{k} \in \left\{\frac{K}{2}, \frac{K+2}{2}, \frac{K+4}{2}\right\}.$$

Using Mathematica, we find that near $\beta = 1$,

$$\Delta\left(\frac{K}{2}, \frac{K+2}{2}, \beta\right) = (\beta-1)^2 + \mathcal{O}\left(|\beta-1|^3\right)$$
$$\Delta\left(\frac{K+2}{2}, \frac{K+4}{2}, \beta\right) = -(\beta-1)^2 + \mathcal{O}\left(|\beta-1|^3\right).$$

It follows that near and below $\beta = 1$ we have $\frac{K}{2}, \frac{K+4}{2} \prec \frac{K+2}{2}$.
For item 1(c), using Mathematica, we find that $\beta \to 0^+$ implies $k^* \to 1$, so that for small positive β, $\hat{k} \in \{1, 2\}$. Also, $\beta \to 0^+ \implies \psi \to 0$, so $\hat{k} \neq 0$. Using Mathematica, we find that near $\beta = 0$,

$$\Delta(1, 2, \beta) = (2 - 4(1+K)^{-1})(\beta-1)^2 + \mathcal{O}\left(|\beta-1|^3\right), \tag{A10}$$

so that $\hat{k} = 2$ for $K \geq 2$. When $K = 1$ we again find $\Delta = 0$, so that $\hat{k} = 1$.
Case 2: $-1 < \beta < 0$. We establish item 2(a) by direct analysis, and noting that it suffices to study the behavior of k^*. Noting that $-1 < \psi < 0$, we compute

$$\log \psi^2 = 2\log(-\psi) = \log\left(\frac{\beta}{\beta-1}\right) + \log\left(1 - \beta^{K+1}\right) - \log(1+K) \to -\infty \tag{A11}$$

$$K^{-1}\log \psi^2 = K^{-1}\log\left(\frac{\beta}{\beta-1}\right) + K^{-1}\log\left(1 - \beta^{K+1}\right) - K^{-1}\log(1+K) \to 0 \tag{A12}$$

Since $k^* = \log \psi^2 / \log \beta^2$ and $\log \beta^2 < 0$ we see that by equation (A11) $k^* \to \infty$, and that by equation (A12) ${}^{k^*}\!/_K \to 0$.

Turning to item 2(b), using Mathematica we find that $\beta \to 0^-$ implies $k^* \to 1$, and since $\beta \in (0, 1)$, we know that $\psi < 0$ so that $\hat{k}$ is odd. It follows that for β is near and below 0 we have $\hat{k} \in \{1, 2\}$. The expansion (A10) then shows that

$\hat{k} = 3$ for $K \geq 2$. Also as before, $K = 1$ implies $\Delta = 0$, so that $\hat{k} = 1$. Finally, for item 2(c), we find using we find that $\beta \to -1^+$ implies $k^* \to \infty$, and the result follows.

Case 3: $\beta < -1$. We establish item 3(a) by direct analysis. First, observe that $\beta < -1$ implies

$$|\psi(K, \beta)| = \left(\frac{\beta}{\beta - 1}\right)\left(\frac{(\beta^2)^{\frac{K+1}{2}} + (-1)^{K+1}}{K+1}\right)$$

By L'Hopital's rule, the function $f(x) = (2\alpha)^{-1}(x^\alpha + \beta)$ diverges to infinity as $\alpha \to \infty$ for $x > 1$ and for any $\beta \in \mathbb{R}$, which shows that $|\psi(K, \beta)| \to \infty$ as $K \to \infty$. It follows that $\log \psi^2 \to \infty$, and thus k^* and $\hat{k}$ go to infinity as $K \to \infty$. Next note

$$\frac{k^*}{K} = (\log(-\beta)^{-1}\left(K^{-1}\log(\beta-1)^{-1}\beta + K^{-1}\log\left((\beta^2)^{\frac{K+1}{2}} + (-1)^{K+1}\right) - K^{-1}\log(K+1)\right).$$

It follows that

$$\lim_{K\to\infty}\frac{k^*}{K} = \lim_{K\to\infty}(K\log(-\beta))^{-1}\log\left((\beta^2)^{\frac{K+1}{2}} + (-1)^{K+1}\right). \tag{A13}$$

Let $g(x) = \alpha^{-1}\log(x^{\frac{\alpha-1}{2}} + \beta)$, for $\beta \in \mathbb{R}$ and $x > 1$. Then

$$\lim_{\alpha\to\infty} g(x) = \lim_{\alpha\to\infty}\frac{x^{\frac{\alpha-1}{2}}\log(x)}{2\left(x^{\frac{\alpha-1}{2}} + \beta\right)} = \log(x)/2.$$

It follows that

$$K^{-1}\log\left((\beta^2)^{\frac{K+1}{2}} + (-1)^{K+1}\right) \to \log(\beta^2)/2 = \log(-\beta),$$

which, when combined with (A13), yields the result.

Turning now to item 3(b), note that if K is odd then $\psi \to 0$, so that $\hat{k} \to 0$. If K is even then $\psi \to -(K+1)^{-1} \in (0, 1)$, so that $\hat{k} \to 1$. Finally, for item 3(c), using Mathematica, we find that $\beta \to -\infty$ implies $k^* \to K+1$. By Lemma A.6 we know

$$\lim_{\beta\to-\infty}\hat{k} \in \{K-1, K+1, K+3\}.$$

Again using Mathematica we find that if $K \geq 2$ then

$$\lim_{\beta\to-\infty}\Delta(K-1, K+1) = \lim_{\beta\to-\infty}\Delta(K+1, K+3) - \infty,$$

so that eventually $K+1, K+3 \prec K-1$. If $K = 1$, then $\Delta(K-1, K+1) = 0$ and so by our tie-breaker, $\hat{k} = 0$. ■

A2 Simulated dynamics of the unified model

To illustrate how convergence is achieved under different specifications of the unified dynamics, we consider a variety of special cases operating under a range of feedback parameters β. In this section, without loss of generality, we set γ

at zero, so that $\bar{y} = 0$ (equivalently, the dynamics for y and a may be viewed as in deviation form). We take the parametric form of the rate function for the replicator dynamics to be given by $r(x) = {}^2\!/\!_\pi \tan^{-1}(\alpha x)$, with $\alpha > 0$. Finally, all simulations are initialized with $a_0 = 1$ and $\omega_{k0} = {}^1\!/\!_4$ for $k = 0, 1, 2, 3$.

Figure A1: Simulated dynamics with positive feedback

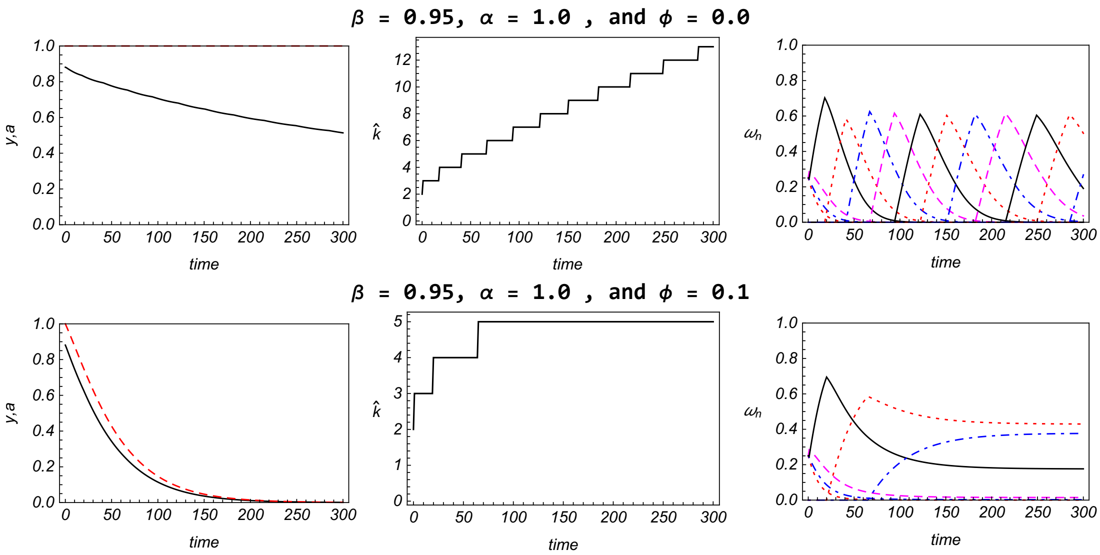

Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). In the left panels the *solid black* curves denote y and in the bottom left panel the *dashed red curve* identifies a. In the right panels $\omega_{n0} = {}^1\!/\!_4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *red dotted, blue dash-dot, dashed magenta* and *solid black*, respectively.

We start with with the stable positive feedback case $0 < \beta < 1$: see Figure A1, where $\beta = 0.95$ and $\alpha = 1$. Upper row corresponds to replicator dynamics ($\phi = 0$) and bottom row to unified dynamics ($\phi = 0.1$): we omit results associated with adaptive dynamics as they simply show monotonic convergence of a and y to $\bar{y}$.

Under replicator dynamics, y exhibits monotone convergence to $\bar{y}$, as the weight distribution shifts to higher k-level forecasts. The upper-right panel provides the dynamics of agents' weights. The time paths for weights $\omega_{n0} = {}^1\!/\!_4$, $n = 0, 1, 2, 3$, are distinguished by plot-style: *red dotted, blue dash-dot, dashed magenta* and *solid black*, respectively. As the replicator adds higher forecast levels, the associated paths are graphically identified in an analogous fashion by repeating the styles mod four. Under replicator dynamics, lower-level forecasts gradually fall out of favor and are replaced by higher-level forecasts.

Under unified dynamics, convergence is now much faster, and also faster than the adaptive dynamics case. The optimal k appears to stall out at $\hat{k} = 5$ because, as the estimate $a_t \to 0$, higher-level forecasts provide limited to no additional value.

We now turn to the negative feedback case, with $-1 < \beta < 0$. The results associated with adaptive dynamics are unexceptional. Figure A2 provides the results for $\beta = -0.5$. Under replicator dynamics, the behavior of y is non-monotonic: the upper-left panel, shows oscillatory convergence of y induced by the negative feedback. The behavior of $\hat{k}$ reflects these oscillations: when y crosses zero, $\hat{k}$ rises sharply to drive down (in magnitude) the optimal forecast $\beta^{\hat{k}}$.

Figure A2: Simulated dynamics with negative feedback.

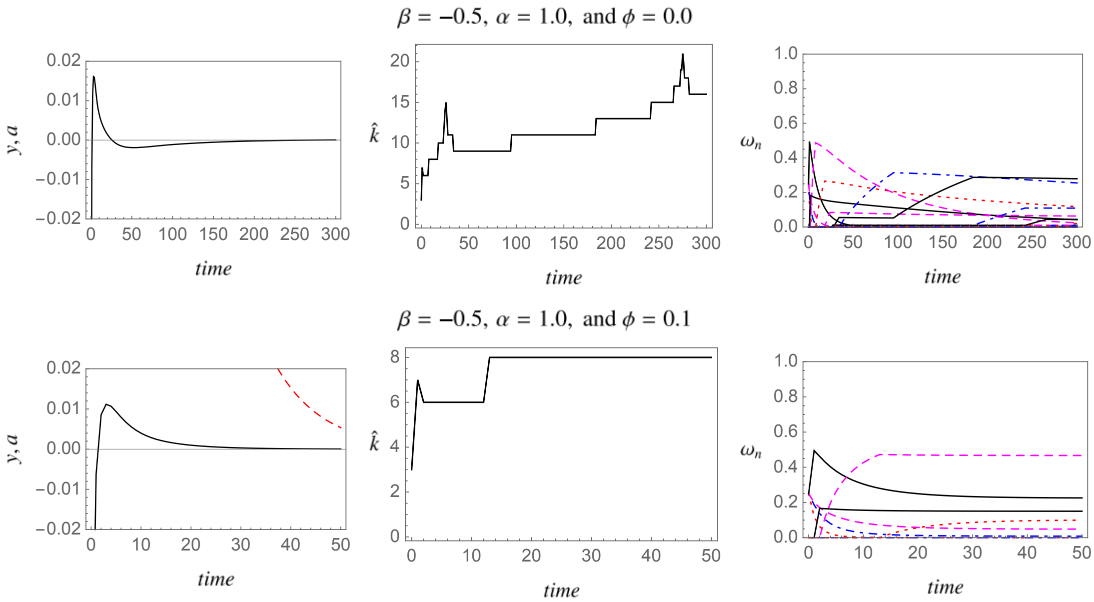

Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *red dotted, blue dash-dot, dashed magenta* and *solid black*, respectively.

By Theorem 2, $\hat{k} \to \infty$. However, unlike the positive feedback case, here this convergence is not monotone. Figure A2 also gives the results for unified dynamics. Because adaptive dynamics drives level-0 forecasts to zero there is faster convergence, with weaker oscillatory behavior, than under the replicator.

Figure A3: Simulated dynamics with large negative feedback.

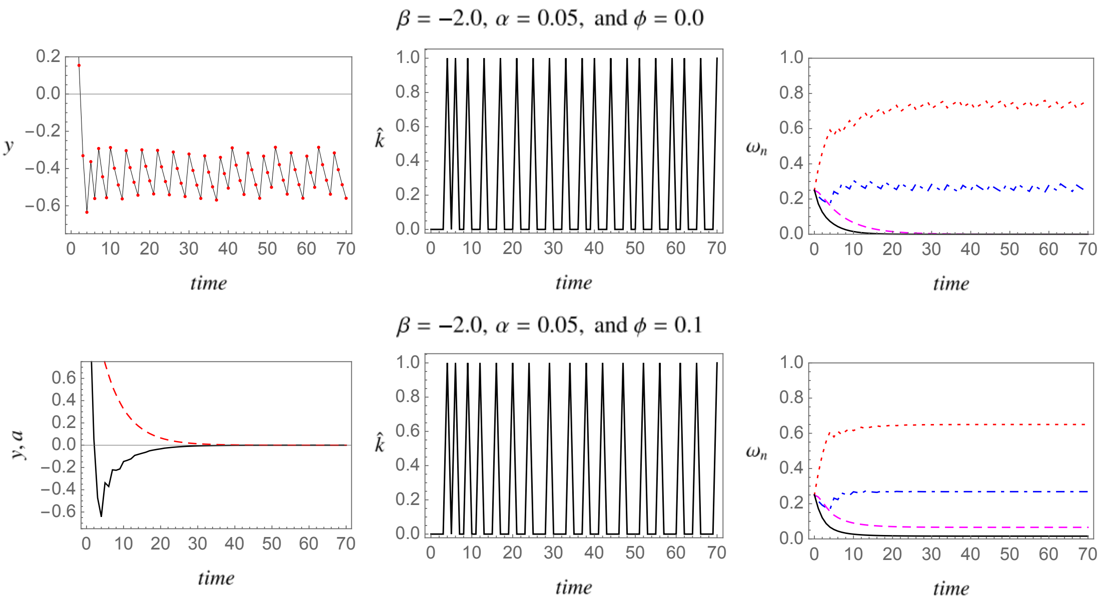

Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *red dotted, blue dash-dot, dashed magenta* and *solid black*, respectively.

Finally, we turn to the case in which $\beta < -1$. We remark that, in this case, $\bar{y}$ is not stable under eductive learning as shown in (n.d.): if all agents are fully rational and have common knowledge of the structure they are unable to coordinate on the REE. However, as indicated by Corollary 1, when $\beta < -1$ the REE is stable under adaptive dynamics provided the gain is sufficiently small.

In the replicator-only case, the dynamics can be unstable or can exhibit complex behavior. For example, the top panel of Figure A3 provides a simulation with $\beta = -2.0$ and $\alpha = 0.05$. Note that $\hat{k}$ oscillates between 0 and 1, which drives ω_{nt} to zero for $n \geq 2$. The evolution of y appears to converge to an 11-cycle, which, we observe, is not centered at zero.[39] The bottom row of Figure A3 exhibits the corresponding simulation with unified dynamics. The addition of adaptive dynamics pushes level-0 expectations towards zero, which when combined with replicator dynamics leads to rapid convergence to the REE.

A3 Additional Experimental Results

Figure A4 shows the average price observed across all treatments relative to the REE price. Figure A5 shows the individual price predictions for all individuals with outliers indicated by X's. The individual forecasts illustrate both the diversity and uniformity that can occur depending on the expectational feedback in the market. As predicted by the simulations shown in Section 4.3, all the $|\beta| < 1$ cases show convergence to the REE initially and after the announcements, whereas both convergence and non-convergence is observed when $\beta < -1$.

We observed more outliers in individual predictions in this study than were observed, for example, in Bao and Duffy (2016). However, we also have more than double the participants. Some outliers are easily explained as "fat finger" errors where an extra zero is added to a forecast. Others reflect participants with a penchant for anarchy who consistently typed in nonsensical forecasts. In fact, we identify two anarchists who repeatedly typed in the highest price permitted just to see what would happen. One of these anarchists actually provided a nice natural experiment within our laboratory experiment, which we discuss in detail in Section A3.1 below.

When classifying individual forecasts without cutoffs, we chose to not classify 35 out of the 18,367 forecasts from our analysis (5 of which occurred in announcement rounds out of 517 observations in total).[40] Nearly half of the total outliers forecasts were submitted by just 3 (out of the 372) participants in the study. The outlier predictions on average were for a price of 391, which is nearly 200 larger than any plausible price in any treatment. If these outliers were classified as level-k, then most are classified as an REE prediction (e.g. in a positive feedback treatment when level-k forecasts converge from below the REE price and the outlier is above the REE price) or a level-0 prediction (e.g. when convergence starts from above an REE price and level-k deductions are closer to the REE price), which is clearly not in keeping with what the classification is attempting to achieve.[41]

[39] We find numerically that there are at least two stable 11-cycles.

[40] The odd number of observations is due to two markets that did not complete the experiment. One market in a T2×A1 treatment ended early when a participant withdrew from the experiment. The other was a T2×A2 treatment that ended a short time after the announcement when a student kicked a power cord knocking out two computers with players in the same market. The data up to that point was saved, but there was no way to let the students pick up where they left off.

[41] Inclusion of these outliers actually makes some of our results stronger. For example, with

Figure A4: Average market price relative to REE

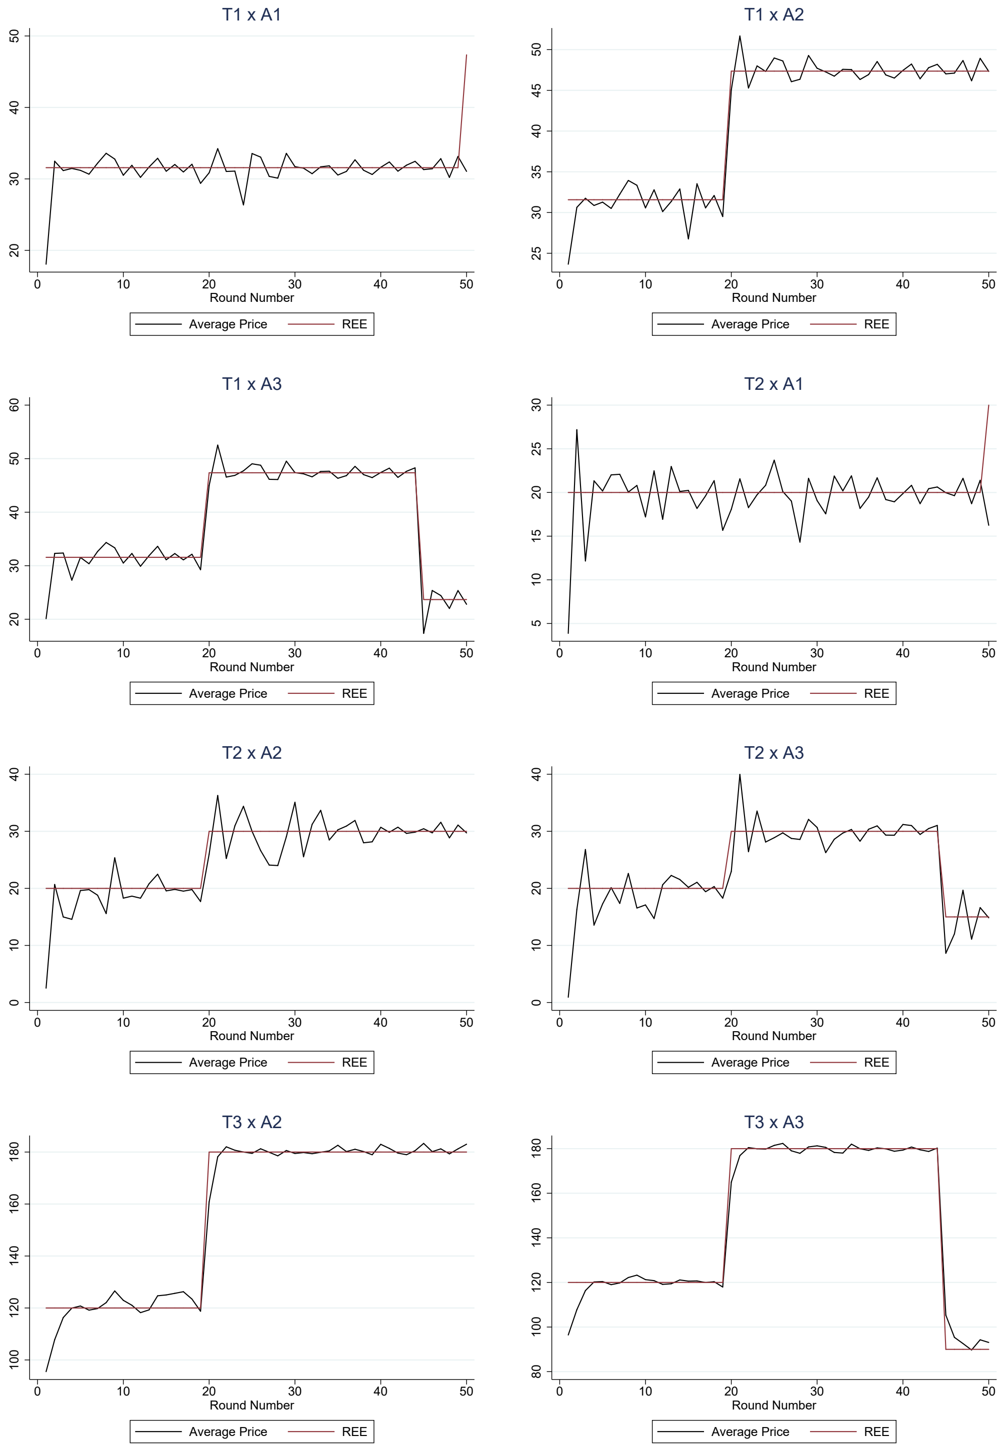

Figure A5: Individual participant predictions

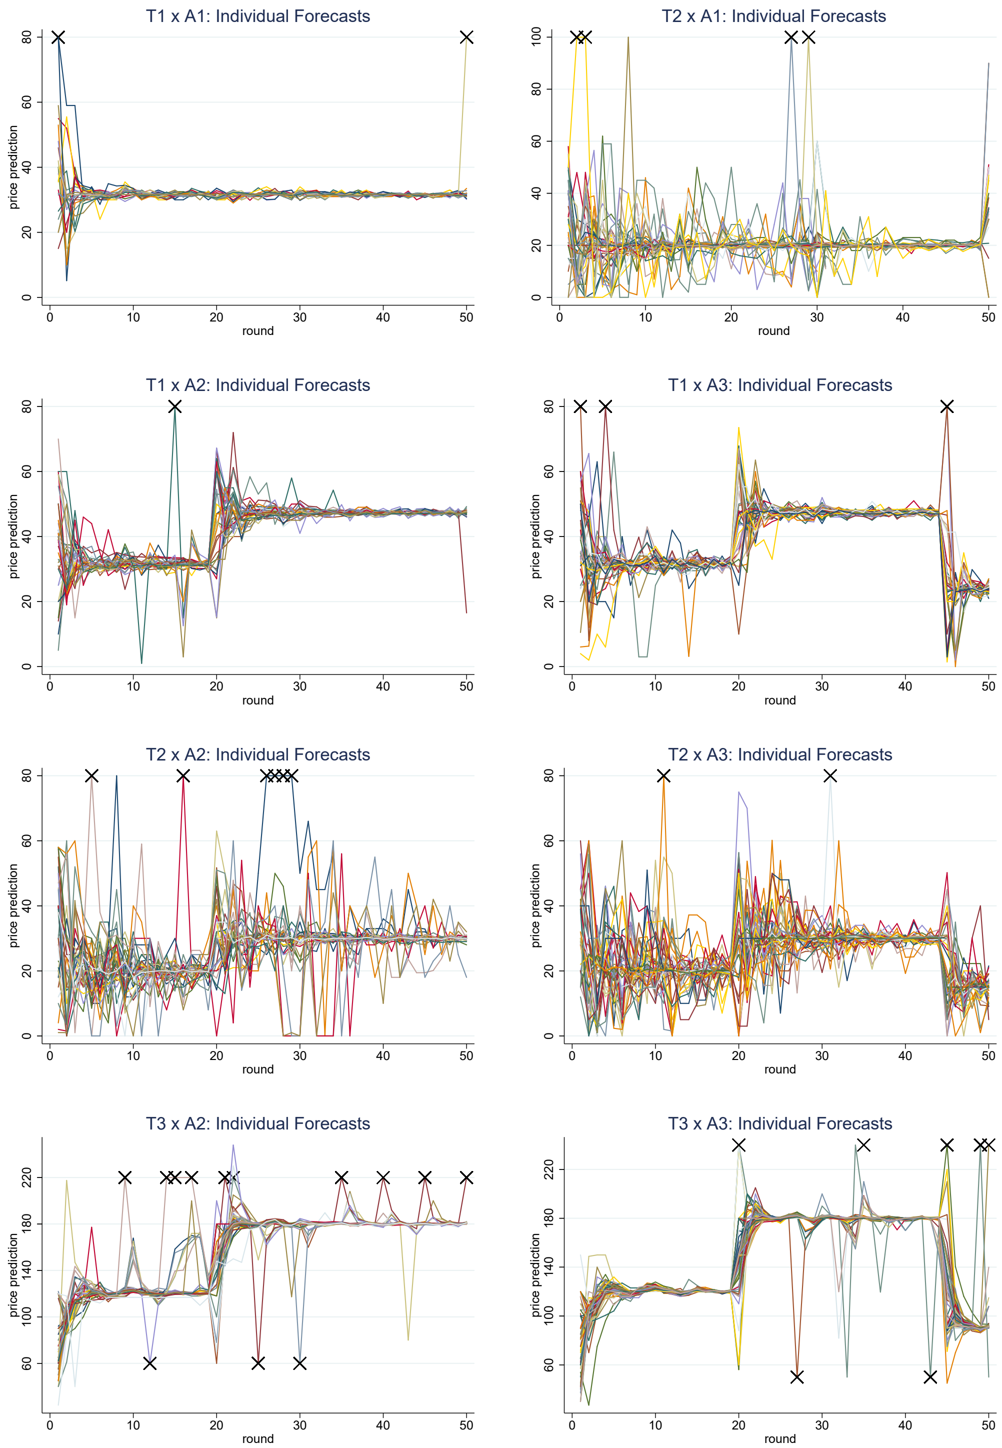

Notes: The 'X's denote forecasts that are larger than the top axis shown in the graph. The maximum value the program would allow a participant to predict is 500.

Table A7 provides an overall breakdown of the data, including the outliers, to provide a sense of how far away most forecasts are from the model predictions. The table shows three measures of the root squared difference between a subjects submitted forecast and the nearest level-k model implied forecast, where the level-k forecasts are constructed using the standard assumptions given in Section 6 in the main text. The root mean squared error/difference (RMSE) for the classifications are quite large. This is almost entirely due to outliers and a minority group of the submitted forecasts. The root median squared error/difference (RMedSE) shows that the majority of forecasts are with one unit of a level-k forecast overall and within 4 units in announcement rounds. The final statistic reported in the table is the 70^{th} percentile of root squared differences. This statistic is chosen because we found that approximately 70% of participants chose a level-k forecast in an announcement round when we use a cutoff value of ± 4.5 for pooled data (see Table 3 in the main text). The column illustrates a treatment-by-treatment breakdown of that classification.

Table A7: Classification of predictions using counterfactual forecast rules

	All observations			Announcement periods only		
Treatment	RMSE	RMedSE	70^{th} Pctl	RMSE	RMedSE	70^{th} Pctl
T1 x A1	14.92	0.31	0.57	73.61	1.00	2.35
T1 x A2	10.53	0.36	0.64	7.78	1.49	4.90
T1 x A3	9.73	0.37	0.78	28.97	1.70	4.35
T2 x A1	9.59	0.30	0.73	7.29	4.00	5.00
T2 x A2	21.78	0.36	1.05	7.34	3.00	5.56
T2 x A3	3.97	0.50	1.32	5.39	3.00	5.00
T3 x A1	22.53	0.50	1.00	12.04	2.00	9.07
T3 x A2	14.72	0.44	1.00	28.07	2.01	5.00

Notes: This table shows how well laboratory participants' forecasts can be classified using a counterfactual forecast. For each subject we construct Level-0, 1, 2, 3, and REE forecasts based on the observed market data available to participants at each point in time. We calculate the difference between this forecast and the observed forecast submitted by the participant. We classify the subject as Level-0, 1, 2, 3, or REE based on which comparison yields the lowest squared error. The table reports the root mean (RMSE), median (RMedSE), and 70^{th} percentile of the squared difference between the submitted forecast and the nearest counterfactual forecast. The 70^{th} percentile is shown because we were able to classify 70% of forecasts in announcement periods using a ±4.5 cutoff when the data is pooled.

respect to the result reported in Table 4 in the main text, the forecast errors generated by some of these outliers move the results in favor of the unified model.

Figure A6: Comparing the unified model to experimental data

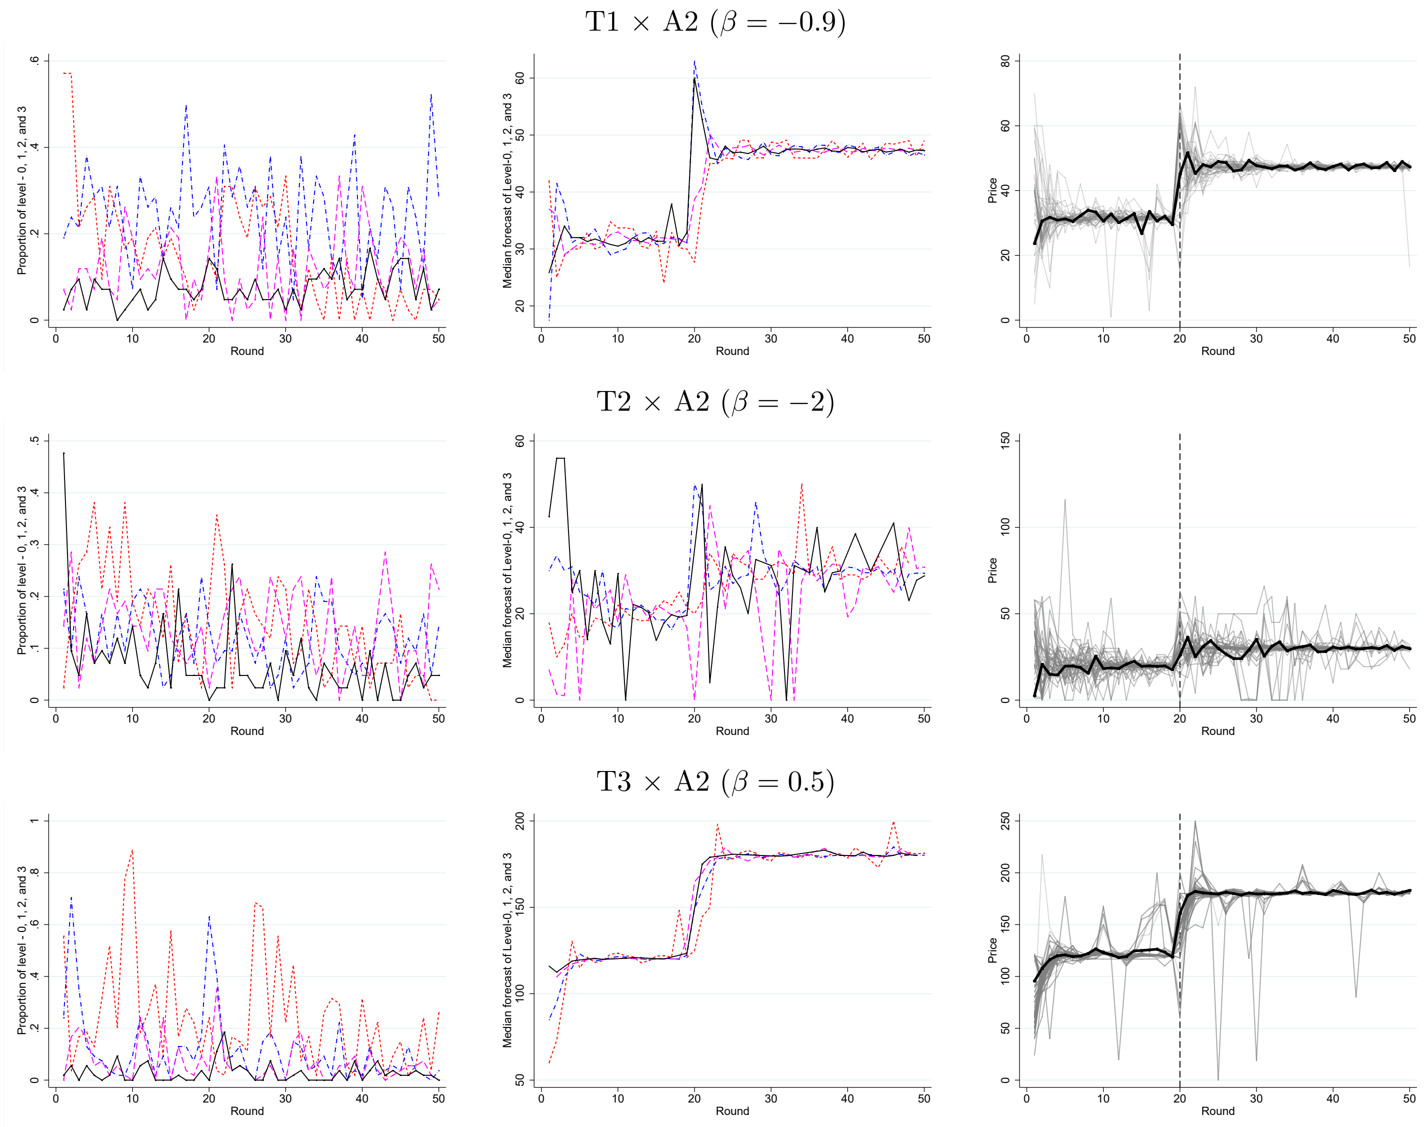

Notes: Survey participants' forecasts are classified as Level-0, 1, 2, 3, or consistent with the REE forecast by comparing to the model implied forecasts. The time path of observed ω_n for $n = 0, 1, 2, 3$ are distinguished by plot-style: *red dotted, blue dash-dot, magenta dash* and *black solid*, respectively. The corresponding median forecasts, $E_{t-1}y_t^k$, of the participants use the same style format. The final column shows average market prices observed (solid black) laid over all individual forecasts. We omitted some outliers from the the final column of figures, which are shown in Figure A5 for clarity. The omitted forecasts are included in the calculations in the first two columns.

Figure A6 provides the same data breakdown for A2 treatments that we provided for A3 treatments in Figure 3 in the main text. We find similar results here. We identify heterogeneous forecasts that display level-k depths of reasoning in announcement rounds with median individual and mean market dynamics closely matching what was predicted in Section 4.3 in the main text.

There are some additional points of interest in Figure A6 worthy of comment. First, results from our simulations could be interpreted as suggesting that the relative proportions of level-k agents would converge over time: see, e.g. Figure A10. This conclusion is in contrast with the first column of Figure A6. However, the failure to observe convergence in the experiment is, for a variety of reasons, to be expected. Most salient is that the simulation is non-stochastic, modeling a continuum of agents, whereas the experiment included explicit stochasticity and only a small number of agents in each market, and lasted fewer rounds.

Second, in the T3 treatment we observed one market that had a significant departures in price from the REE after a period of convergence to the REE. You can see the individual forecasts in the third graph on the right of the last row. There is a group of individual forecasts that rise for many periods prior to the announcement in period 20. This market is what causes the spike in the median level-0 forecast that can be seen in the middle figure on the bottom row of Figure A6. The cause of this divergence is an anarchist player. This player's actions provide a nice case study for the unified model. For the five players who are attempting to play the game normally, the market has both large unobserved shocks and announced shocks.

A4 An Anarchist Anecdote

Figure A7 and A8 provide some detail on this anarchist's market. The first graph in the top left of Figure A7 shows the market price and the individual forecasts of the market participants. The anarchist is shown in red. The market converged to the REE by period 7. The anarchist then decided in period 9 to enters a price of 500, which was the largest price that the program would allow. The next figure shows the result. The price increased and a significant forecast error was realized by all other market participants. The anarchist struck again in round 14 and this time repeatedly enter a price of 500 for four consecutive rounds (ending in round 17). As before, there is a significant forecast error realized by all other players in the period the anarchists defects. However, the players quickly adapt to this unexplained rise in the price and the average forecast error falls over the next four periods. Importantly, we see all players switching to a forecast that lines up well with an adaptive forecast, consistent with the assumptions of the unified model. When the anarchist switches strategy in round 18, another large forecast error is generated, which causes yet another clear change in the strategy choices among the other participants.

The final figure in the top row of A7 layers onto the individual expectations the implied level-0, 1, 2, and 3 forecasts using our standard assumptions from Section 6. The bottom row of figures in A7 zooms in on the period of interest and plots the implied path of a single level-k forecast on each graph for clarity. It is immediately apparent that each large forecast error generates a shift in behavior by the non-anarchist players. Each shift in behavior is well-captured by one of the level-k deductions.

Figure A7: An Anarchist Anecdote

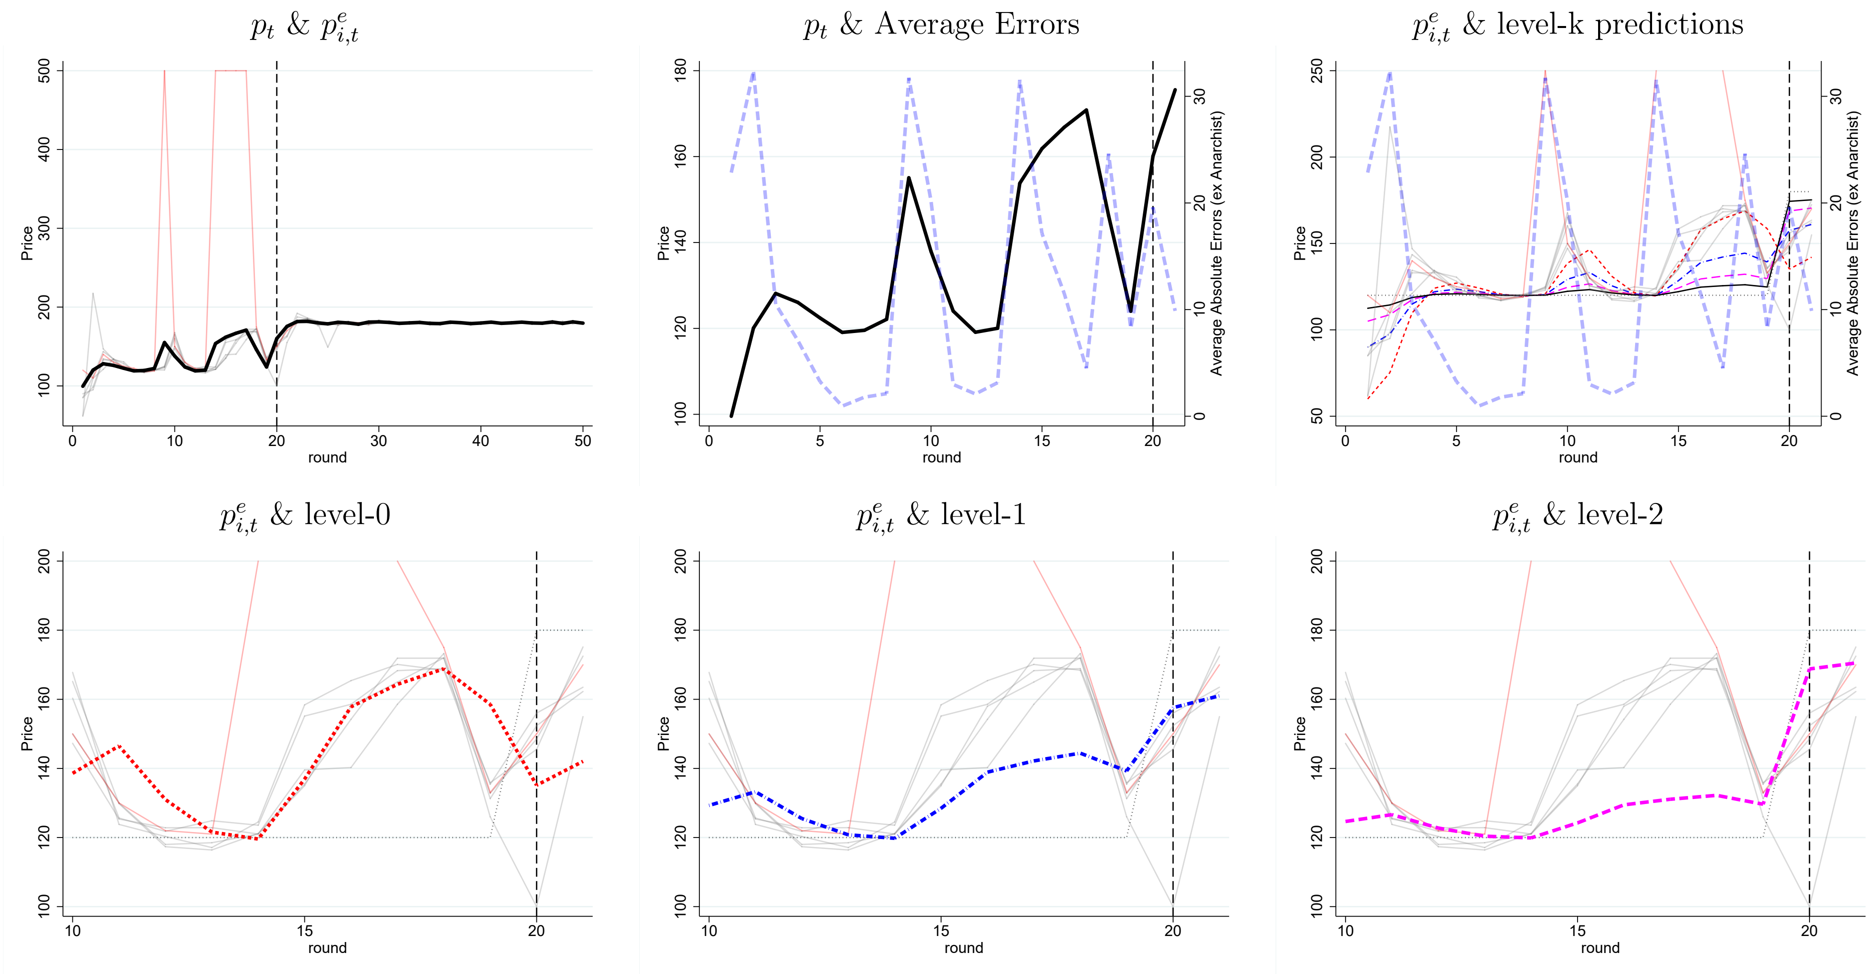

Notes: These figures show data from a T3×A2 treatment experimental market, where one player decided to actively sabotage the market. The anarchist's forecasts are shown in red. The time path of the implied level-0, 1, 2, and 3 forecasts are distinguished by plot-style: *red dotted, blue dash-dot, magenta dash* and *black solid*, respectively. The REE forecast is *black dotted*. The market price is the *solid thick black line.*

To see this, start by looking at period 10. Recall that there is no information that the participants have to suggest why the price suddenly moved in period 9. All participants trend follow in period 10 and revise their forecasts up. But the anarchist reverses course and provides a reasonable forecast in period 10, this generates another sizable forecast error. For period 11, the other participants switch strategies again. They appear to revise up their depth of reasoning and predict that market price will again fall. Both level-1 and level-2 predictions, which are based on the average price for rounds 9 and 10, explain nearly all the variation in forecasts chosen in this period. This switch by participants to a higher level strategy in period 11 generates a low forecast error and the subjects appear to maintain these strategies in the subsequent periods leading the market to converge.

When the anarchists strikes again in period 14, the remaining participants are quick to revise their depth of reasoning down to level-0. Forecast errors fall when switching to this strategy so they maintain the level-0 strategy. When the anarchists stop choosing 500 and reverts to choosing a normal strategy, another large forecast error is realized by the other market participants. This leads to a change in strategy in the next round. The revised strategies observed in the next round all sit on, or between, the implied level-1 and level-2 strategies (see bottom row of plots in Figure A7).

The chaos of this market is distinct from most other markets we observed. This raises the question of what the participants will do in an announcement round after the market has been so unpredictable. It appears that they mostly respond in accordance with the unified model. Five out the six forecasts for the announcement round sit between the level-1 forecast using our standard definition and a level-1 forecast where the level-0 assumptions is $p = 120$, which is the steady state price prior to the announced change.

Figure A8 zooms in even further on just rounds 20 and 21 and classifies the individual forecasts types using the method described in Section 6 in the main text. Between the two rounds of play, those subjects whose forecasts were closest to the actual price, i.e. experienced the smallest errors, stick with the level-1 forecast. Those subjects who experience larger errors clearly revise up their depth of reasoning, where a revision to level-2 corresponds to what would have been the best forecast to play in round 20 given what occurred. This behavior is consistent with the assumptions that underlie the replicator dynamic's reflective process that we assume for the unified model.

A5 Robustness: Level-0 forecast definition

To classify the types of forecasting strategies that participants use, we must assume a shared level-0 forecast. Our baseline assumption is that level-0 is a two-round moving average of past prices. To demonstrate that our results are robust to this assumption, we conduct two exercises. First, we replicate the results in Table 3 and 4 of the main text using a four round moving average as the shared level-0 forecast. Second, we study how overall classification of types and of the level-0 type changes when we assume last periods price as the level-0 forecast, a two-period moving average, a four-period moving average, or three different cosntant gain specifications.

Table A10 replicates Table 3. The number of people we classify as level-k

Figure A8: An Anarchist Anecdote Announcement Round

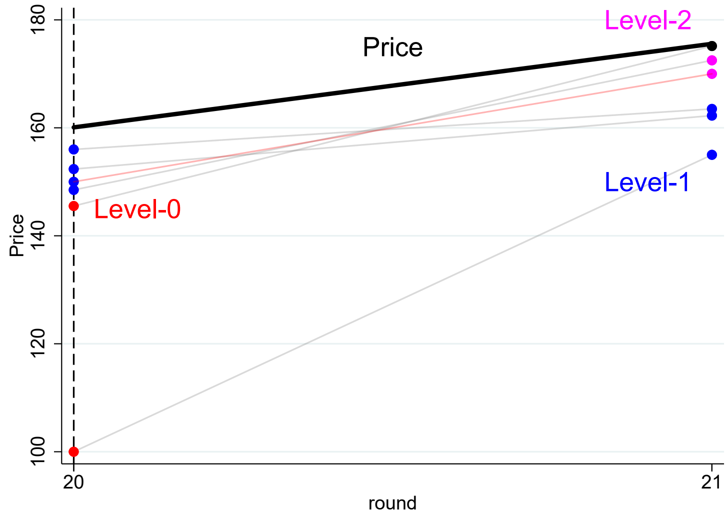

Notes: Individual classified price forecasts in a T3×A2 treatment. The classifications are made by comparing the forecasts to different implied level-k forecasts. The closest implied forecast type determines the classification (see Section 6).

reasoners increases slightly under this definition overall. The regression estimates are mostly unchanged. We retain statistical significance for the hypothesis test conducted on announcement rounds with a comparable F-stat obtained to the original specification.

Table A9 replicates Table 4 for the four-round average level-0 assumption. The results are slightly stronger on all categories relative the previous definition.

Table A10 shows the classification results for the ±3 cut off for different level-0 assumptions. In general, the proportion of subjects that we classify as level-k forecasters of any type increases as we consider level-0 forecasts with longer averages or weighted averages of past observed prices.

A6 Oscillating deductions with strategic substitutes

Figures A9 shows an additioanl example of a market with clear oscillating behavior and Figures A10 shows one market with more monotonic behavior. The top row of plots in each figure summarizes the market dynamics in each case with the first plot showing the market price and all of the individual forecasts, the second plot showing the individual forecasts plotted against the level-0,1,2, and REE forecasts, and the third plot showing the classification of each type along with just the level-0 forecast.[42] The latter two plots are zoomed in around the announcement period. The bottom six plots of each figure show each individual's forecasts from each market classified by level-k type period-by-period compared to just the level-0 forecast. These plots illustrate the evolution of an individual's forecasts over time relative to the principle reference point for level-k deductions: the level-0 forecast.

[42]We do not plot the level-3 forecast in the middle figure because it makes the graph harder to interpret by requiring a larger scale of the y-axis. For the markets we show, no one chooses it in the announcement round. This of course is not true in general. We observe people choosing exactly level-3 deductions in some markets as can be seen in Figure 3 in the main text.

Table A8: Classifying participant's forecasts as Level-k - Robustness Check

Within ±3 of Level-k in announcement rounds	1	20/50	45
Total Classified	47.3%	65.8%	66.0%
	[33.9% , 57.0%]	[50.6% , 73.6%]	[49.4% , 69.9%]
Level-0	14.8%	7.8%	5.1%
	[11.0% , 15.1%]	[4.31% , 9.48%]	[4.49% , 6.41%]
Level-1	7.3%	25.0%	14.1%
	[6.45% , 8.60%]	[19.3% , 27.6%]	[14.1% , 14.1%]
Level-2	6.5%	5.2%	3.8%
	[1.89% , 6.45%]	[4.31% , 5.75%]	[1.92% , 3.85%]
Level-3	3.2%	3.2%	4.5%
	[1.07% , 1.13%]	[2.58% , 3.74%]	[3.21% , 5.13%]
REE	15.6%	24.7%	38.5%
	[13.4% , 15.6%]	[20.1% , 27.0%]	[25.6% , 40.4%]
N	372	348	156

Hypothesis tests of deliberation time regressions	
H_0 : Level-0 - Level-3 = 0	F(1, 61) =0.75
H_0 : (Level-0 x Ann) - (Level-3 x Ann) = 0	F(1, 61) =**4.33**

Differences in deliberation time (seconds)		
Variable	[1]	[2]
Level-0	**-9.73**	**-1.97**
	(1.113)	(0.672)
Level-1	**-8.61**	-0.16
	(1.144)	(0.629)
Level-2	**-5.50**	-0.24
	(1.312)	(0.991)
Level-3	**-6.67**	-0.33
	(1.332)	(1.054)
Level-0 x Ann	**47.52**	3.58
	(8.623)	(6.065)
Level-1 x Ann	**45.47**	**11.90**
	(4.881)	(4.741)
Level-2 x Ann	**8.61**	11.08
	(9.058)	(8.546)
Level-3 x Ann	**63.82**	**23.02**
	(11.92)	(8.260)
Cons	**41.16**	**112.52**
	(0.526)	(4.227)
Individual FE	yes	yes
Round FE	no	yes
R-squared	0.030	0.253
N	18,367	18,367

Notes: The top left panel reports the proportion of participant's forecasts that fall within ±3 of a Level-k forecast. Proportions for cutoffs of ±1.5 and ±4.5 are shown in brackets. The right panel reports the regression results of identified Level-k individual's deliberation time in all periods and in announcement periods. Standard errors are clustered at the market level and reported in parenthesis below the point estimates. Bolded values indicate statistical significance at the ten percent level. The bottom left panel reports the hypothesis tests for the equality of regression coefficients for regression specification (2). We pool A1 (round 50 announcement) and A2 (round 20 announcement) results because both experiments feature a single and identical announcement.

Figures A9, 5 (in main text), and A10 clearly suggest a willingness by some individuals to oscillate their predictions above and below both the level-0 forecast and the REE price. The oscillations occur despite the experience of the price not oscillating for many periods prior to the announcement. This experience of tranquility combined with how close many of the forecasts are to level-k deductions is at least suggestive evidence that people contemplated oscillations consistent with classic level-k reasoning. Moreover, they took action with money at stake consistent with such deductions.

To further illustrate point, Figures A11 show the same type of analysis applied to a T3 treatment where $\beta = 0.5$. Level-k deductions do not imply oscillations in this case and indeed none are observed. Individual forecasts conform nicely to level-k deductions based on our proposed level-0 forecast. This suggests that people do not abandon level-k deductions in environments with strategy substitutability. Level-k deductions describes forecasting behavior in our experiment when strategic actions are both compliments and substitutes.

Table A9: Revisions and loss - Robustness Check

Treatment	Proportion of changers Between rounds 20 & 21		Ave. abs. prediction error Round 20			Ave. deliberation time (sec) Round 21		
	Revise opt.	No Change	Change	No change	Difference	Change	No change	Difference
T1 x A2 and A3	**0.62**	0.37	18.43	7.08	**11.35**	58.3	47.8	10.48
	[5.75]	(31/84)			[6.11]			[1.27]
T2 x A2 and A3	**0.48**	0.51	24.95	13.16	**11.79**	62.7	56.5	6.16
	[3.20]	(46/90)			[4.41]			[0.69]
T3 x A2 and A3	**0.56**	0.25	26.83	14.84	**11.99**	36.5	35.8	0.73
	[6.24]	(30/119)			[3.31]			[0.12]
	Between rounds 45 & 46		Round 45			Round 46		
T1 x A3	**0.73**	0.64	26.4	5.34	**21.06**	43.3	30.0	**13.3**
	[3.97]	(27/42)			[6.39]			[1.35]
T2 x A3	0.30	0.58	17.42	11.77	**5.65**	36.1	27.8	**8.23**
	[0.47]	(28/48)			[1.48]			[1.32]
T3 x A3	**0.52**	0.15	31.18	21.58	**9.60**	25.4	18.2	**7.16**
	[4.25]	(10/66)			[1.56]			[2.21]

Notes: "Revise opt." is the proportion of people who, conditioning on changing their strategy in period 21(46), changed their strategy to the best counterfactual strategy out of level-0, 1, 2, 3, or the REE in their market, where best is defined as what forecast would have been best in round 20(45). Z-scores for the test of the null hypothesis that subjects switched to one of the five strategies at random are reported in brackets. The next column reports the proportion of participants who we classify as not changing their strategy either between rounds 20 and 21 or between rounds 45 and 46 following announcements in either round 20 or 45, respectively. Counts appear in parentheses below. The remaining columns report the difference in average absolute prediction errors and average deliberation time for subjects classified as changing versus not changing with two-sample t-test statistics reported in brackets. Bolded values represent statistical significance at the ten percent level.

Table A10: Classifying participant's forecasts as Level-k - Robustness Check

	Moving averages			Constant Gain		
Round 20/50	1-period	2-period*	4-period	$\phi = 0.4$	$\phi = 0.3$	$\phi = 0.2$
Level - k	63.8%	64.4%	65.8%	65.2%	66.1%	66.1%
Level - 0	6.6%	6.6%	7.8%	6.9%	7.2%	7.2%

Notes: *Assumption used for level-k classification in the main text. The table reports the proportion of participant's forecasts that fall within ±3 of a Level-k forecast in all treatments with an announcement in period 20 or 50. The level-k forecasts are based on the level-0 assumption denoted in the table.

Figure A9: Example 1: Individual forecasts from experimental market with treatment T2 ($\beta = -2$)

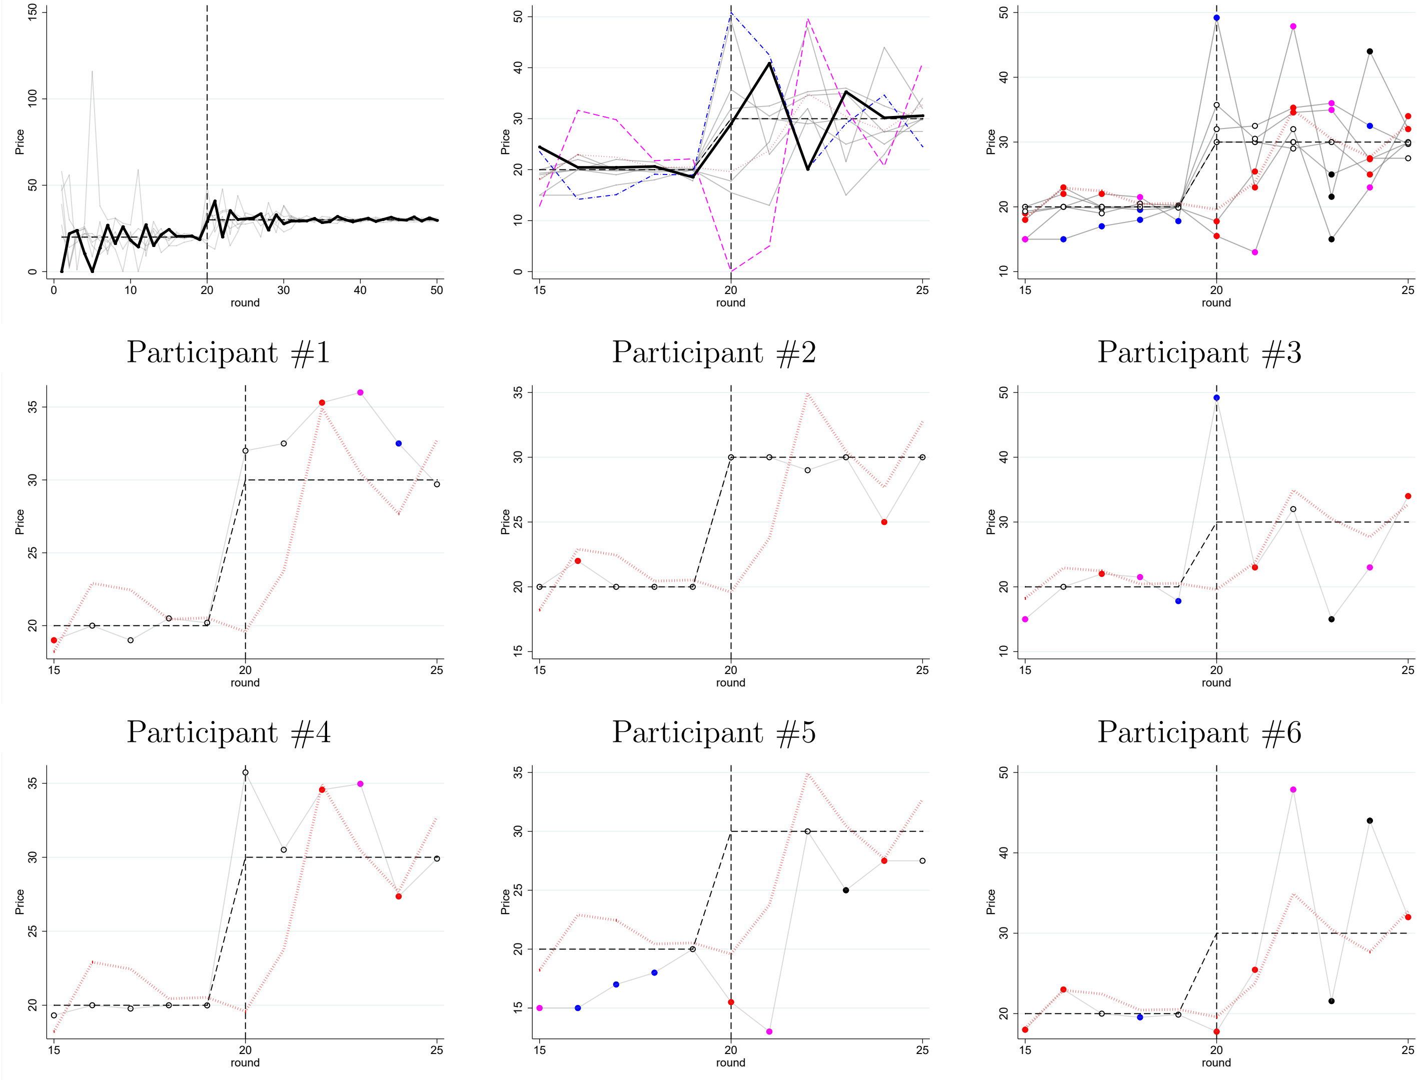

Notes: The first plots shows all individual forecasts and the market price from a single market. The second plot shows the model implied level-k forecasts and the data in a window around the announcement. The remaining figures classify each of the forecasts as a level-k type, which is indicated by the color of the dot. Forecasts that are classified as level-0 are shown in red, level-1 in blue, level-2 in magenta, level-3 in black, and REE as a black circle with a white interior. The dotted red line shows the path of the level-0 forecast from which all level-k deductions are derived. The perfect foresight equilibrium is indicated by the dashed line.

Figure A10: Example 3: Individual forecasts from experimental market with treatment T2 ($\beta = -2$)

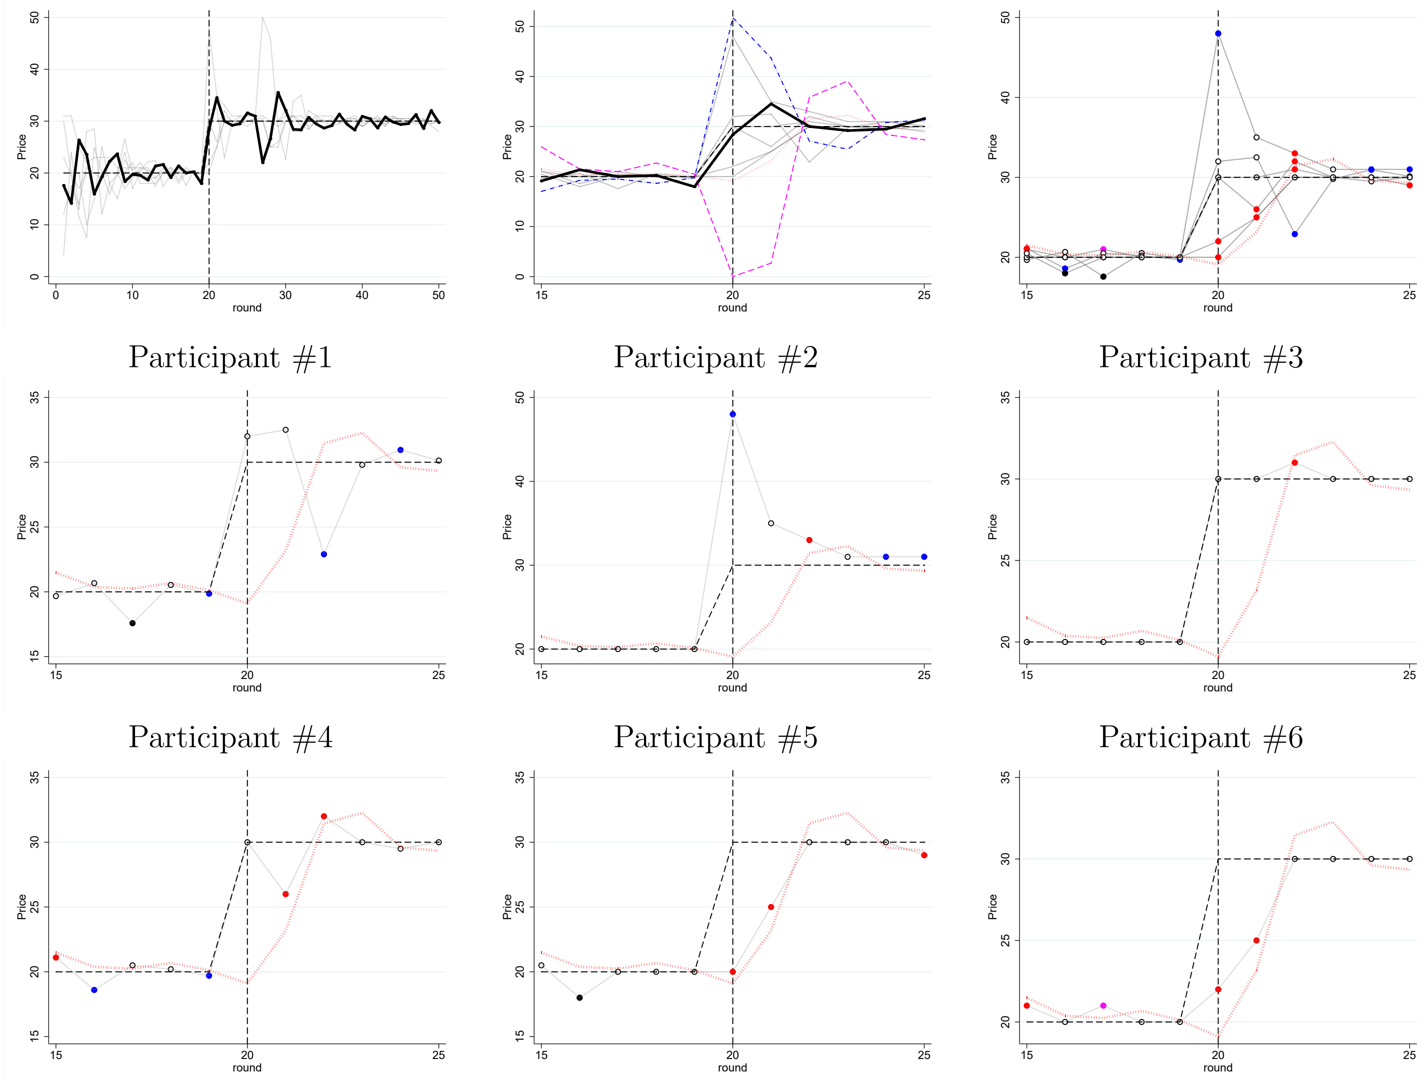

Notes: The first plots shows all individual forecasts and the market price from a single market. The second plot shows the model implied level-k forecasts and the data in a window around the announcement. The remaining figures classify each of the forecasts as a level-k type, which is indicated by the color of the dot. Forecasts that are classified as level-0 are shown in red, level-1 in blue, level-2 in magenta, level-3 in black, and REE as a black circle with a white interior. The dotted red line shows the path of the level-0 forecast from which all level-k deductions are derived. The perfect foresight equilibrium is indicated by the dashed line.

Figure A11: Example 4: Individual forecasts from experimental market with treatment T3 ($\beta = 0.5$)

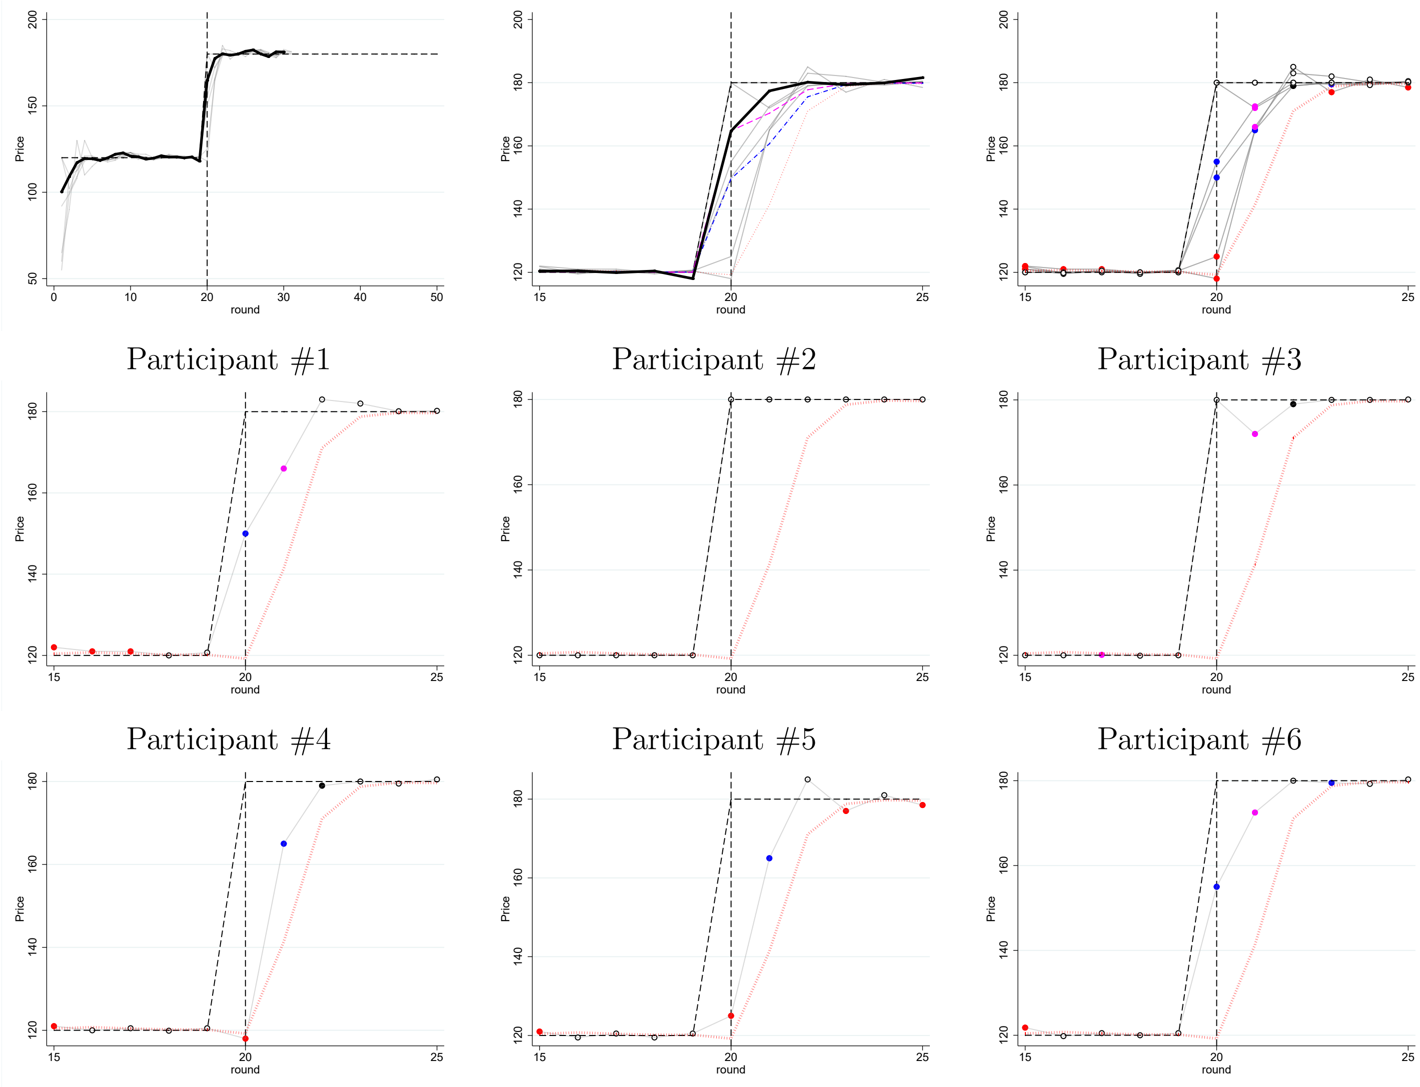

Notes: The first plots shows all individual forecasts and the market price from a single market. The second plot shows the model implied level-k forecasts and the data in a window around the announcement. The remaining figures classify each of the forecasts as a level-k type, which is indicated by the color of the dot. Forecasts that are classified as level-0 are shown in red, level-1 in blue, level-2 in magenta, level-3 in black, and REE as a black circle with a white interior. The dotted red line shows the path of the level-0 forecast from which all level-k deductions are derived. The perfect foresight equilibrium is indicated by the dashed line.

A7 Exit Survey Results

After the experimented ended, subjects completed an exit survey while they waited for their pay envelopes to be prepared. The survey questions aimed to assess what information they used to make their forecasts and what information they thought others used.

A7.1 Exit Survey Questions

1. Please rank the importance of each option below to the formation of your price forecast in each period:

   a. The history of market prices

   b. The market equations

   c. The history of my own price forecasts

   d. The history of my own forecasts errors

   e. My expectation about the average price forecast in the period

2. Please rank the importance of each option below to the formation of your price forecast **following the announcements**:

   a. The history of market prices

   b. The market equations

   c. The history of my own price forecasts

   d. The history of my own forecasts errors

   e. My expectation about the average price forecast in the period

3. Which of the following statements best describes your thinking before making each forecast?

   a. I looked at the past prices and made my best guess based on their recent movements. I never used the equations.

   b. I made a guess about what the average forecast might be based on past prices and then used the equations to determine my own forecast using that guess.

   c. I made a guess about what the average forecast might be and used the equation to work out the price only when I did a poor job of forecasting in the previous round. Otherwise, I just looked at past prices and made my best guess.

   d. I made a guess about what the average forecast might be and used the equation to work out the price only when there was an announced change in the market. Otherwise, I just looked at past prices and made my best guess.

4. Please rank the importance of each option below to **other participants**, which you believe they may have used to make their price forecasts:

    a. The history of market prices

    b. The market equations

    c. The history of their own price forecasts

    d. The history of their own forecasts errors

    e. Their expectation about the average price forecast in the period

5. Please rank the importance of each option below to **other participants**, which you believe they may have used to make their price forecasts **following the announcements:**

    a. The history of market prices

    b. The market equations

    c. The history of their own price forecasts

    d. The history of their own forecasts errors

    e. Their expectation about the average price forecast in the period

6. If you do not feel like the strategy you used was well-captured by the survey questions, then please use this box to explain your strategy

A7.2 Exit Survey Results

Survey questions (1), (2), (4), and (5) used a drop-down menu with options: "very important", "somewhat important", and "did not consider." Table A11 and Table A12 shows the cumulative importance of each factor where "very important" is assigned a zero, "somewhat important" is assigned a one, and "did not consider" a two. Therefore, the lower the value, the more important the information. Consistent with level-k reasoning, we find that on average subjects rated the equations and the forecast of the average expectation as more important to their own forecast than they believed it was to others. This is consistent with a belief that others are less sophisticated. We observe the results on the full sample and when restricting to only people who played a level-k forecast in the announcement periods with the $\pm$ 3 cutoff. The latter consistently rank the equations as important to them than they are to their perceived competitors, which is consistent with the level-k assumption that others players are perceived as less sophisticated.

Figure A12 shows the responses to question 3 separated by treatment. The most common response is (b), which is:

> *I made a guess about what the average forecast might be based on past prices and then used the equations to determine my own forecast using that guess.*

This response is consistent with level-1 behavior.

Table A11: Tabulated survey results for Q1 and Q4

All Responses										
	Past Prices		Equations		Forecast History		Forecast Errors		Exp. Ave. Price	
Treatment	Own	Others	Own	Others	Own	Others	Own	Others	Own	Others
T1×A1	10	9	33	32	24	21	30	32	19	21
T1×A2	11	11	32	30	33	19	38	31	17	19
T1×A3	11	11	31	30	40	31	45	32	21	23
T2×A1	13	8	25	23	31	25	36	30	16	19
T2×A2	13	10	28	28	47	24	50	42	18	30
T2×A3	23	18	36	41	52	35	47	43	25	41
T3×A2	7	7	42	37	50	33	44	42	20	29
T3×A3	16	13	47	59	67	52	63	66	30	33
All	104	87	274	280	344	240	353	318	166	215
Difference	17		-6		104		35		-49	
Info is (__) to me	(less important)		(more important)		(less important)		(less important)		(more important)	
Responses from those identified as level-k in announcement rounds with ±3 cutoff										
	Past Prices		Equations		Forecast History		Forecast Errors		Exp. Ave. Price	
Treatment	Own	Others	Own	Others	Own	Others	Own	Others	Own	Others
T1×A1	6	6	24	22	13	10	17	18	10	13
T1×A2	5	7	10	14	18	12	21	15	7	9
T1×A3	4	3	3	7	8	9	8	8	3	5
T2×A1	7	3	8	11	14	9	13	9	8	9
T2×A2	4	3	2	7	18	12	23	20	6	14
T2×A3	3	4	1	6	6	4	5	2	2	10
T3×A2	0	0	5	6	14	10	14	9	6	9
T3×A3	6	4	10	14	12	12	11	14	4	8
All	35	30	63	87	103	78	112	95	46	77
Difference	5		-24		25		17		-31	
Info is (__) to me	(less important)		(more important)		(less important)		(less important)		(more important)	

Notes: Participants rated each piece of information denoted in the top line as “very important”, “somewhat important”, or “did not consider” when making their “own” forecasts and what they believed was important to “others”. The categories are assigned the following values and summed: “very important” is a assigned a zero, “somewhat important” a one, and “did not consider” as two. Lower totals indicate that the piece of information is more important to a person’s decision.

A8 Quantitative evaluation: additional results

We fit the model to the experimental data at the market level. Table 5 in the main text averages over the individual market outcomes from the same treatments. Table A13 shows the underlying data from each market.

Each model that features heterogeneous types is initialized to the first realized price and to the distribution of level-k types observed in period one for each market. Afterwards, the model makes predictions based solely on the evolution of price, adaptive learning, or the replicator, depending on which model is used. The learning model starts initial beliefs at the average of the individual forecasts in period one. After period one it updates according to the evolution of data implied by the model and beliefs for the chosen gain. The simulated data is compared to experimental data and the mean squared error is calculated.

Each model is optimized individually by searching over a grid of gains $\phi \in [0, 1]$, or replicator parameters $\alpha \in [0, 2]$, or both in the case of the unified model. The optimal coefficients are shown in Table A14. Both the replicator and adaptive learning are required to best fit the data in T1 and T2 treatments. In many of the T3 treatments, however, naive expectations and fixed level-k reasoning is chosen as the best model. This reflects the fact that many markets coverge very quickly

Table A12: Tabulated survey results for Q2 and Q5

All Responses										
	Past Prices		Equations		Forecast History		Forecast Errors		Exp. Ave. Price	
Treatment	Own	Others	Own	Others	Own	Others	Own	Others	Own	Others
T1×A1	13	12	30	29	24	24	29	39	23	18
T1×A2	19	19	24	19	42	35	36	32	23	21
T1×A3	19	19	28	24	40	29	44	38	24	23
T2×A1	11	21	18	23	33	27	40	34	21	20
T2×A2	21	19	21	20	44	24	49	41	20	24
T2×A3	31	35	30	36	52	49	54	45	32	35
T3×A2	25	27	31	25	51	39	54	45	26	28
T3×A3	40	44	38	38	75	64	73	76	28	30
All	179	196	220	214	361	291	379	350	197	199
Difference	-17		6		70		29		-2	
Info is (__) to me	(more important)		(less important)		(less important)		(less important)		(more important)	
Responses from those identified as level-k in announcement rounds with ±3 cutoff										
	Past Prices		Equations		Forecast History		Forecast Errors		Exp. Ave. Price	
Treatment	Own	Others	Own	Others	Own	Others	Own	Others	Own	Others
T1×A1	7	7	22	21	12	15	18	23	13	10
T1×A2	11	9	6	4	22	19	16	19	9	12
T1×A3	4	3	4	6	8	8	8	10	5	5
T2×A1	7	7	7	12	16	8	16	13	9	11
T2×A2	9	8	1	4	19	10	22	18	14	8
T2×A3	4	4	1	5	5	4	4	1	10	7
T3×A2	1	4	7	9	15	9	18	10	9	7
T3×A3	11	11	8	7	19	12	15	17	8	7
All	54	53	56	68	116	85	117	111	77	67
Difference	1		-12		31		6		10	
Info is (__) to me	(less important)		(more important)		(less important)		(less important)		(less important)	

Notes: Participants rated each piece of information denoted in the top line as "very important", "somewhat important", or "did not consider" when making their "own" forecasts and what they believed was important to "others". The categories are assigned the following values and summed: "very important" is a assigned a zero, "somewhat important" a one, and "did not consider" as two. Lower totals indicate that the piece of information is more important to a person's decision.

to steady state, but not as quickly as RE implies. This is also reflected in the results for the adaptive learning case were a naive model is found to best fit the data for all markets. In subsequent exploration, which is not shown here, we have found that a $\phi > 1$ plus level-k reasoning is preferred. That is consistent with a trend following behavior similar to what many other positive feedback experiments have found.

A8.1 Unified dynamics in the NK model

Consider the following RE forward model:

$$y_t = \gamma_t + \beta E_t y_{t+1}, \tag{A14}$$

where $y \in \mathbb{R}^n$, $\gamma_t \in \{\bar{\gamma}_1, \ldots, \bar{\gamma}_M\}$, $\bar{\gamma}_i \in \mathbb{R}^n$, and γ_t has transition matrix P. The analog of (A14) in our unified framework is

$$y_t = \gamma_t + \beta \sum_k \omega_{kt} \cdot E_t^k y_{t+1}, \tag{A15}$$

where $E_t^k y_{t+1}$ is the level-k forecast of y_{t+1} made in period t.

Table A13: MSE between experimental data and competing models

Treatment	REE	Unified Model		Fixed Level-k		Replicator only		Adaptive learning	
T1 × A3	MSE	MSE	Rel. REE	MSE	Rel. REE	MSE	Rel. REE	MSE	Rel. REE
Market 1	6.67	3.88	0.58	11.59	1.74	4.29	0.64	20.77	3.11
Market 2	8.66	4.17	0.48	13.34	1.54	4.17	0.48	27.38	3.16
Market 3	24.18	22.34	0.92	25.07	1.04	22.34	0.92	39.76	1.64
Market 4	14.47	3.01	0.21	3.87	0.27	3.51	0.24	29.48	2.04
Market 5	14.66	2.24	0.15	12.81	0.87	13.96	0.95	7.85	0.54
Market 6	18.20	4.11	0.23	16.94	0.93	17.74	0.97	17.23	0.95
Market 7	5.22	1.89	0.36	2.94	0.56	2.57	0.49	13.32	2.55
Average	13.15	5.95	0.45	12.37	0.94	9.80	0.74	22.26	1.69
T2 × A3									
Market 1	66.42	57.92	0.87	126.77	1.91	76.33	1.15	86.21	1.30
Market 2	25.31	20.44	0.81	154.20	6.09	34.67	1.37	34.06	1.35
Market 3	58.01	76.90	1.33	873.05	15.05	94.48	1.63	77.46	1.34
Market 4	48.70	40.98	0.84	779.52	16.01	75.16	1.54	65.73	1.35
Market 5	23.36	37.13	1.59	80.20	3.43	44.65	1.91	42.05	1.80
Market 6	44.84	51.04	1.14	569.37	12.70	69.59	1.55	68.70	1.53
Market 7	67.28	52.25	0.78	671.08	9.98	75.55	1.12	47.06	0.70
Market 8	80.64	50.42	0.63	127.50	1.58	97.45	1.21	85.84	1.06
Average	51.82	48.38	0.93	422.71	8.16	70.98	1.37	63.39	1.22
T3 × A3									
Market 1	22.49	1.80	0.08	1.80	0.08	35.16	1.56	36.23	1.61
Market 2	31.16	14.70	0.47	16.33	0.52	46.91	1.51	39.09	1.25
Market 3	37.48	17.64	0.47	17.64	0.47	42.97	1.15	38.23	1.02
Market 4	31.37	13.70	0.44	13.70	0.44	31.78	1.01	48.27	1.54
Market 5	12.52	4.34	0.35	4.34	0.35	28.54	2.28	48.90	3.91
Market 6	28.67	33.11	1.15	35.38	1.23	60.66	2.12	75.75	2.64
Market 7	45.41	23.82	0.52	27.08	0.60	61.14	1.35	46.89	1.03
Market 8	44.70	19.56	0.44	22.89	0.51	50.28	1.12	48.38	1.08
Market 9	92.75	76.53	0.83	76.53	0.83	104.95	1.13	106.02	1.14
Market 10	31.71	3.46	0.11	3.46	0.11	40.40	1.27	28.27	0.89
Market 11	30.64	9.45	0.31	9.45	0.31	41.05	1.34	39.53	1.29
Average	37.17	19.83	0.53	20.78	0.56	49.44	1.33	50.51	1.36

Notes: Mean square error (MSE) of five simulated models of aggregate price dynamics compared to experimental market price data. "Rel. REE" reports the MSE of the a model relative to REE MSE, i.e., Model MSE/REE MSE. Models are fit by doing a grid search over values $\alpha \in [0, 2]$ and $\phi \in [0, 1]$.

Table A14: Parameter estimates of competiting models

Treatment	Unified Model		Fixed Level-k		Replicator only		Adaptive learning
T1 × A3	α	ϕ	α	ϕ	α	ϕ	ϕ
Market 1	0.225	0.725	-	0.475	0.025	-	0.425
Market 2	0.015	0.000	-	0.550	0.015	-	0.500
Market 3	0.008	0.000	-	0.325	0.008	-	0.450
Market 4	0.200	0.775	-	1.000	0.005	-	0.475
Market 5	0.175	0.725	-	1.000	0.000	-	0.550
Market 6	0.150	0.750	-	1.000	0.000	-	0.575
Market 7	0.300	0.800	-	0.725	0.010	-	0.525
T2 × A3							
Market 1	0.005	0.100	-	0.000	0.005	-	0.325
Market 2	0.010	0.050	-	0.475	0.010	-	0.325
Market 3	0.010	0.050	-	0.125	0.010	-	0.400
Market 4	0.005	0.200	-	0.150	0.010	-	0.500
Market 5	0.015	0.025	-	0.175	0.010	-	0.300
Market 6	0.010	0.025	-	0.000	0.010	-	0.325
Market 7	0.005	0.175	-	0.150	0.010	-	0.525
Market 8	0.025	0.425	-	0.400	0.010	-	0.375
T3 × A3							
Market 1	0.000	1.000	-	1.000	0.175	-	1.000
Market 2	0.600	0.725	-	1.000	0.200	-	1.000
Market 3	0.000	1.000	-	1.000	0.175	-	1.000
Market 4	0.000	0.950	-	0.950	0.025	-	1.000
Market 5	0.000	1.000	-	1.000	0.075	-	1.000
Market 6	0.600	0.725	-	1.000	0.375	-	1.000
Market 7	0.600	0.725	-	1.000	0.350	-	1.000
Market 8	0.100	0.725	-	1.000	0.375	-	1.000
Market 9	0.000	1.000	-	1.000	0.200	-	1.000
Market 10	0.000	1.000	-	1.000	0.200	-	1.000
Market 11	0.000	1.000	-	1.000	0.225	-	1.000

Notes: Paremeter estimates of the competiting models. Models are fit by doing a grid search over values $\alpha \in [0, 2]$ and $\phi \in [0, 1]$. The Fixed level-k model assumes an adaptive level-0 forecast.

Figure A12: Exit survey question 3 responses

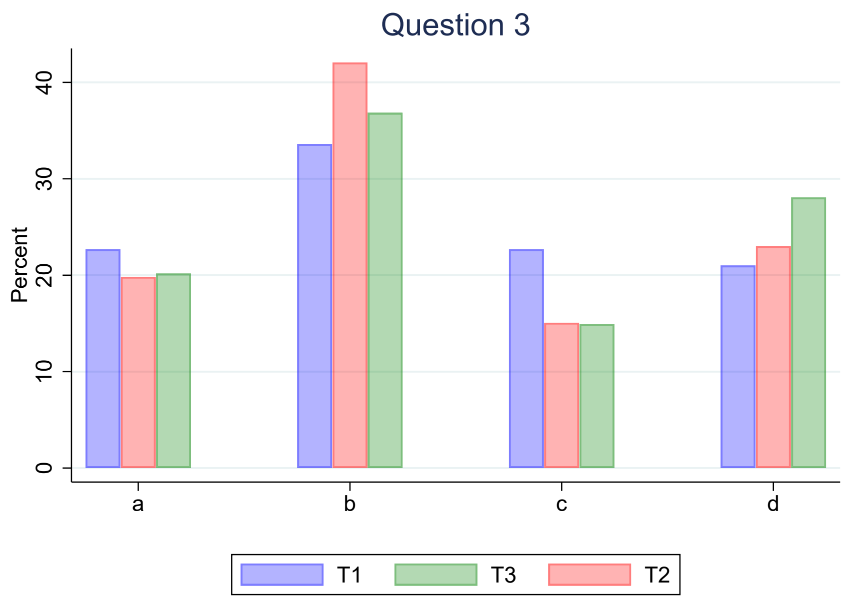

Notes: This figure shows the response to question 3 from exit survey separated by treatment type.

Denote by $\mathcal{E}_i^k y_t$ the level-k forecast of y_t given $\gamma_{t-1} = \bar{\gamma}_i$. Note that the subscript on $\mathcal{E}$ identifies the current state. Importantly, this operator is always a one-period-ahead forecast. Define $\mathcal{E}_i^k y_t$ recursively:

$$\begin{aligned} \mathcal{E}_i^k y_{t+1} &\equiv E_t^k \left(\gamma_{t+1} + \beta E_{t+1}^{k-1} y_{t+2} \middle| \gamma_t = \bar{\gamma}_i \right) \\ &= \sum_j P_{ij} \bar{\gamma}_j + \beta \sum_j P_{ij} \cdot \mathcal{E}_j^{k-1} y_{t+2}. \end{aligned} \tag{A16}$$

Writing $\mathcal{E}^k y_{t+1}$ as the column vector with i^{th} entry as $\mathcal{E}_i^k y_{t+1}$, we obtain

$$\begin{aligned} \mathcal{E}^k y_{t+1} &= (P \otimes I_n)\bar{\gamma} + (P \otimes \beta)\mathcal{E}^{k-1} y_{t+2} \\ &= (I_{Mn} - P \otimes \beta)^{-1} \left(I_{Mn} - (P \otimes \beta)^k \right) (P \otimes I_n)\bar{\gamma} + (P \otimes \beta)^k \mathcal{E}^0 y_{t+k+1}. \end{aligned} \tag{A17}$$

Equation (A17) extends the univariate and two-state case shown in the main text. What is clear though from this more general setting is the connection between the k^{th}-level of reasoning and how forward-looking an agent is. Note the k^{th} exponent in the first term and that the whole term reflects a finite sum of k elements. Each higher level deduction is essentially is akin to contemplating a longer possible duration with the addition of one additional element.

Connecting (A17) to (A15) is simplified by defining $Y_{it} = y_t \big| \gamma_t = \bar{\gamma}_i$, and letting Y_t be the vector of vectors Y_{it} for $i = 1, \ldots, M$. Thus Y_t is the vector of state-contingent values of y_t (the time-subscript on y continues to have relevance because the value of y is history dependent via level-0 forecasts). Exploiting this notation, we may write

$$Y_t = \bar{\gamma} + (I_M \otimes \beta) \sum_k \mathcal{E}^k y_{t+1}. \tag{A18}$$

Given level-0 beliefs, equation (A18) determines y_t for all possible realizations of γ_t.

To complete the model we must determine the evolution of $\{\mathcal{E}^0 y_{t+s}\}_{s\geq 1}$ over time t. We assume agents use CGL coupled with the anticipated utility assumption that $\mathcal{E}_j^0 y_{t+s} = \mathcal{E}_j^0 y_{t+1} \equiv a_{jt-1}$, where this second equality aligns the notation used here with notation from the paper. We have

$$a_{jt} = a_{jt-1} + \chi_j(\gamma_t) \cdot \phi \cdot (y_t - a_{jt-1}). \tag{A19}$$

The function χ_j controls updating of beliefs depending on the realization of the state. If χ_j is the state-j indicator then this algorithm is state-contingent CGL.

Application to Bilbiie's NK model

Recall our laboratory NK model:

$$x_t = E_t x_{t+1} - \sigma^{-1}(i_t - E_t \pi_{t+1} - r_t) \tag{A20}$$

$$\pi_y = \rho E_t \pi_{t+1} + \kappa y_t \tag{A21}$$

Bilbiie assumes r_t is Markov, $r_0 = r_S < 0$, and that it transitions to the absorbing state $r_N > 0$ with probability $\delta \geq 0$. The interest rate rule sets $i_t = 0$ when $r_t = r_S$; otherwise, it transitions to $i_t = r_N$ with probability $\nu \geq 0$.

To map this model into the general theory developed above, we proceed as follows: writing $y_t = (x_t, \pi_t)$ (of course, vectors are always columns), let $\xi_t \in \{S, F, N\}$ be the underlying state, so that $r_t = r(\xi_t)$ and $i_t = i(\xi_t)$. Then $\xi_0 = S$ and the transition for ξ_t is

$$P = \begin{pmatrix} 1-\delta & \delta(1-\nu) & \delta\nu \\ 0 & 1-\nu & \nu \\ 0 & 0 & 1 \end{pmatrix}. \tag{A22}$$

Set

$$\beta = \begin{pmatrix} 1 & \sigma^{-1} \\ \kappa & \rho + \kappa\sigma^{-1} \end{pmatrix} \text{ and } \gamma_t = \begin{pmatrix} \sigma^{-1}(r_t - i_t) \\ \kappa\sigma^{-1}(r_t - i_t) \end{pmatrix}. \tag{A23}$$

Then (A20) – (A21), together with the interest rate rule, can be written as (A14). In more detail,

$$r_t(\xi) = \begin{cases} r_S & \text{if } \xi = S \\ r_N & \text{if } \xi = F \\ r_N & \text{if } \xi = N \end{cases} \text{ and } i_t(\xi) = \begin{cases} 0 & \text{if } \xi = S \\ 0 & \text{if } \xi = F \\ r_N & \text{if } \xi = N \end{cases}$$

Thus

$$\gamma_t(\xi) = \begin{cases} \sigma^{-1}(r_S, \kappa r_S)' & \text{if } \xi = S \\ \sigma^{-1}(r_N, \kappa r_N)' & \text{if } \xi = F \\ (0,0)' & \text{if } \xi = N \end{cases}$$

Finally, to incorporate our assumptions about level-0 forecasts, we assume that level-0 forecasts in state S are the same as in state F, and level-0 forecasts in state N are rational, i.e. zero. Operationally, we have that $E_t^0\left(y_{t+1}\middle|\xi_t=\xi\right)=a_{\xi t}$, where

$$a_{\xi t}=a_{\xi t-1}+\chi_\xi(\xi_t)\cdot\phi\cdot(y_{t-1}-a_{\xi t-1}),$$

and

$$\chi_N(\xi)=0,\ \ \chi_S(\xi)=\begin{cases}1 & \text{if } \xi=S\\ 0 & \text{else}\end{cases},\ \ \chi_F(\xi)=\begin{cases}1 & \text{if } \xi=S \text{ or } F\\ 0 & \text{else}\end{cases},$$

and $a_{S,0}=a_{F,0}$, and $a_{N,0}=0$.

The REE of the model is obtained by backward induction. Let y_S and y_F be the output-gap/inflation pairs associated with γ_S and γ_F respectively. Since $y=0$ when $\gamma=\gamma_N$ we can solve for y_S and y_F using backward induction. We have

$$y_F=\gamma_F+P_{22}\beta y_F,\ \text{thus}\ y_F=\left(I_2-P_{22}\beta\right)^{-1}\gamma_F \tag{A24}$$

$$y_S=\gamma_S+P_{11}\beta y_S+P_{12}\beta y_F,\ \text{thus}\ y_S=\left(I_2-P_{11}\beta\right)^{-1}\gamma_S+P_{12}\beta\left(I_2-P_{22}\beta\right)^{-1}\gamma_F \tag{A25}$$

A9 Experiment materials

This section provides the instructions and tutorial information that were provided to laboratory subjects.

Negative feedback case

Computer based tutorial:

- What is your role?

  Your role is to act as an expert forecaster advising firms that produce widgets.

- What makes you an expert in this market?

  You will have access to information about the demand and supply of widgets to the market. You will also have a bit of training before making paid forecasts.

- What is a widget?

  Widgets are a perishable commodity like bananas or grapes. They are perishable in the sense that they can only be consumed in the period they are produced. They cannot be stored for consumption in future periods. The widgets that each firm produces are all the same and there are many firms in the market. Therefore, the individual firms do not set the price at which they sell their widgets but must sell widgets at the market price.

- Why do the firms need to forecast the price?

  A firm must commit to the number of widgets it will produce in the coming period before knowing the price. Therefore, the firms need to have a forecast of the price to know how many to produce.

- How am I paid?

Your compensation for each forecast is based on the accuracy of the forecast. The payoff for each forecast is given by the following formula:

$$\text{payment} = 0.50 - 0.03\,(p - \text{your price forecast})^2$$

where p is the actual market price, and 0.50 and 0.03 are measured in cents. If your forecast is off by more than 4, you will receive \$0.00 for your forecast. Therefore, you will receive \$0.50 for a perfect forecast, where p=your price forecast, and potentially \$0.00 for a very poor forecast. You will be paid to make 50 forecasts in total.

In addition, you will be paid a \$5 show-up fee for participating. You may quit the experiment at any time, for any reason, and retain this \$5 payment.

- The Demand for Widgets:

  The total demand for widgets in a period is downward sloping. This means that the lower the price is the greater the demand for widgets. In precise terms, the demand is given by

$$q = A - Bp$$

  where q is the quantity demanded, and p is the current price in the market. The equation for demand and the values for A and B will be given to you at the beginning of the experiment. The values may also change during the experiment. The equation, the values of A and B, and any changes to these values will be told to all participants at the same time.

- The Supply of Widgets:

  The firms in the market all face the same costs for producing widgets. The supply of widgets by each firm, therefore, only depends on their forecast for price next period. The total supply of widgets to the market depends on the average price forecast from all firms.

  The total amount of widgets supplied to the market by all firms is given by

$$q = D \times \text{average price forecast}$$

  where D is a positive number, which will be given to you and all other forecasters in the market at the start of the experiment. Just like with demand, D may change during the experiment and the changes will be announced.

- Prices and Expected Prices:

  Once all participants have chosen their expected price, the average expected price determines total supply. Since quantity demanded depends on price, equating supply and demand determines the price. Consequently the actual market price depends on average expected price. In fact there is a negative relationship between price and expected price. In other words, when the average forecast for the price is high, the actual price is low and vice versa.

- Why does this occur?

  It occurs because a high average expected price causes widget producers to increase their production of widgets. The increase in production results in more widgets supplied to the market. More supply of widgets means that the price of each widget will be lower. The opposite occurs when the average expected price is low. In this case, the widget producers will supply fewer widgets to the market, which results in a high price.

By equating supply and demand,

$$A - Bp = D \times \text{average price forecast}$$

we can arrive at the precise relationship for price and expected price

$$p = \frac{A}{B} - \frac{D}{B} \times \text{average price forecast}$$

Note that expected price is negatively related to price. If expected price is high, then the actual price is low and vice versa

- A bit of randomness:

  Finally, like in real markets, we allow for the possibility that unforeseen and unpredictable things may happen that affect price. We add this to the game by adding a **small** amount of noise to price such that

$$p = \frac{A}{B} - \frac{D}{B} \times \text{average price forecast} + \textit{noise}.$$

  The *noise* term is chosen at random in each period and is not predictable. Its value is not given to any participant in the market. The size of each realisation is **small**. The average value of the *noise* over the course of the experiment is zero and each realisation of it is independent from any other realization. In other words, the *noise* term may take a positive or a negative value in any given period, but overall, the size and number of positive and negative realisations will be approximately equal and cancel each other out over time.

Positive feedback case

Computer based tutorial:

- What is your role?

  Your role is to act as an expert forecaster advising firms that sell widgets.

- What makes you an expert in this market?

  You will have access to information about the demand and supply of widgets to the market. You will also have a bit of training before making paid forecasts.

- What is a widget?

  Widgets are a perishable commodity like bananas or grapes. They are perishable in the sense that they can only be consumed in the period they are produced. They cannot be stored for consumption in future periods. The widgets are all the same and there are many firms that sell in the market. Therefore, the individual firms do not set the price at which they sell their widgets but must sell widgets at the market price.

- Why do the firms need to forecast the price?

  Widgets are considered by many to be a luxury good, in part because they cannot be stored. In fact, when the price of widgets goes up, the demand for widgets tends to go up as well as many consider expensive widgets a status symbol. Therefore, how many widgets a firm should produce to meet demand depends on the expected price in the market that day. Each firm has an advisor like you that provides price forecasts. If the average price forecast is high, then firms will want to supply many widgets and the actual price will be high. If the average price forecast is low, then the firms will supply fewer widgets and the actual price will be low.

- How am I paid?

  Your compensation for each forecast is based on the accuracy of the forecast. The payoff for each forecast is given by the following formula:

  $$\text{payment} = 0.50 - 0.03\,(p - \text{your price forecast})^2$$

  where p is the actual market price, and 0.50 and 0.03 are measured in cents. If your forecast is off by more than 4, you will receive \$0.00 for your forecast. Therefore, you will receive \$0.50 for a perfect forecast, where p=your price forecast, and potentially \$0.00 for a very poor forecast. You will be paid to make 50 forecasts in total.

  In addition, you will be paid a \$5 show-up fee for participating. You may quit the experiment at any time, for any reason, and retain this \$5 payment.

- The Demand for Widgets:

  The total demand for widgets in a period is upward sloping. This means that the higher the price, the greater the demand for widgets. In precise terms, the demand is given by

  $$q = A + Bp$$

  where q is the quantity demanded, and p is the current price in the market. The equation for demand and the values for A and B will be given to you at the beginning of the experiment. The values may also change during the experiment. The equation, the values of A and B, and any changes to these values will be told to all participants at the same time.

- The Supply of Widgets:

  The firms in the market all face the same costs for producing widgets. The supply of widgets by each firm, therefore, only depends on their advisor's forecast for price next period. The total supply of widgets to the market depends on the average price forecast from all firms.
  The total amount of widgets supplied to the market by all firms is given by

  $$q = C + D \times \text{average price forecast}$$

  where C and D are positive numbers, which will be given to you and all other forecasters in the market at the start of the experiment. Just like with demand, C and D may change during the experiment and the changes will be announced.

- Prices and Expected Prices:

  Once all advisors have chosen their expected price, the average expected price determines total supply. In each period, a central market-maker then sets the final price so that demand equals the quantity supplied. Consequently, the actual market price depends on the average expected price. In fact, there is a positive relationship between price and expected price. In other words, when the average forecast for the price is high, the actual price is high and vice versa.

- Why does this occur?

  It occurs because a high average expected price causes widget producers to increase their production of widgets. The higher the price, the higher the actual demand for widgets due the fact they are a status symbol. The opposite occurs when the average expected price is low. In this case, low prices will results in low demand as widgets appear to be less of a luxury good. By equating supply and demand,

$$A + Bp = C + D \times \text{average price forecast}$$

we can arrive at the precise relationship for the price and the expected price

$$p = \frac{C - A}{B} + \frac{D}{B} \times \text{average price forecast}$$

where we will assume that $C > A$. Note that the expected price is positively related to price. If the expected price is high, then the actual price is high and vice versa

- A bit of randomness:

  Finally, like in real markets, we allow for the possibility that unforeseen and unpredictable things may happen that affect price. We add this to the game by adding a **small** amount of noise to price such that

$$p = \frac{A}{B} - \frac{D}{B} \times \text{average price forecast} + noise.$$

  The *noise* term is chosen at random in each period and is not predictable. Its value is not given to any participant in the market. The size of each realisation is **small**. The average value of the *noise* over the course of the experiment is zero and each realisation of it is independent from any other realization. In other words, the *noise* term may take a positive or a negative value in any given period, but overall, the size and number of positive and negative realisations will be approximately equal and cancel each other out over time.

Paper instructions:

Widget Game Instruction Summary:

- Your job is to forecast the price of a widget next period
- Demand for widgets is determined by the market price
    - $q = A - B p$
- The total supply of widgets to the market is determined by the **average** of all price forecasts submitted to the market
    - $q = D \times$ average price forecast
- Combining supply and demand, we have the **key formula** that determines price in the market
    - $P = A/B - D/B \times$ average expected price + *noise*
        - Recall that *noise* is small and **on average equal to zero**
- **An Example:** $A = 120$, $B = 2$, $D = 1$, and *noise* $= 0$, what is price if the average price forecast is 42?
    - $p = 60 - ½ \times$ average price forecast
    - $P = 60 - ½ \times 42 = 60 - 21 = 39$
- You are paid based on **accuracy of your forecast** according to the following formula
    - Payment = 0.50 – 0.03 $(p - your\ price\ forecast)^2$
        - A perfect forecast in a round earns 50 cents
        - A very poor forecast results in 0.00
- **KEY POINT: The market has negative feedback. Therefore, if the average price forecast is high, the market price will be low. And, if the average price forecast is low, then the market price will be high.**
- **Your Notes:**
    - –
    - –
    - –
    - –
    - –
    - –

Widget Game Rules

- You may withdraw from the experiment at any time for any reason
- You may take notes on this paper or the scratch paper provided
- Feel free to do any calculations you wish on the scratch paper provided
- **Do not exit the web browser**
- Do not open new tabs in the web browser
- **Please turn your phone off during the experiment**
- Do not speak with the people around you

Widget Game Instruction Summary:

- Your job is to forecast the price of a widget next period
- Demand for widgets is determined by the market price
    - $q = A + B\,p$
- The total supply of widgets to the market is determined by the **average** of all price forecasts submitted to the market
    - $q = C + D\ x$ average price forecast
- Combining supply and demand, we have the **key formula** that determines price in the market
    - $P = (C - A)/B + D/B\ x$ average expected price + *noise*
        - Recall that *noise* is small and **on average equal to zero**
- **An Example:** $A = 0$, $B = 2$, $C = 60$, $D = 1$, and $noise = 0$, what is price if the average price forecast is 42?
    - $p = 30 + \frac{1}{2}\ x$ average price forecast
    - $P = 30 + \frac{1}{2} \times 42 = 30 + 21 = 51$
- You are paid based on **accuracy of your forecast** according to the following formula
    - Payment = 0.50 – 0.03 $(p - your\ price\ forecast)^2$
        - A perfect forecast in a round earns 50 cents
        - A very poor forecast results in 0.00
- **KEY POINT: The market has positive feedback. Therefore, if the average price forecast is high, the market price will be high. And, if the average price forecast is low, then the market price will be low.**
- **Your Notes:**
    - –
    - –
    - –
    - –
    - –
    - –

Widget Game Rules

- You may withdraw from the experiment at any time for any reason
- You may take notes on this paper or the scratch paper provided
- Feel free to do any calculations you wish on the scratch paper provided
- **Do not exit the web browser**
- Do not open new tabs in the web browser
- **Please turn your phone off during the experiment**
- Do not speak with the people around you