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 27, 2022

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. Macroeconomic applications to the cost of disinflation and to forward guidance illustrate the explanatory power of the unified model.

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. The strong assumptions on agents' knowledge and cognitive abilities that RE imposes 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. Indeed, there are numerous empirical puzzles associated with structural change when RE is imposed. To resolve these puzzles, many modelers are turning to boundedly rational alternatives, including adaptive learning (e.g. Evans, Honkapohja and Mitra (2009) and Gibbs and Kulish (2017)), level-k reasoning (e.g. Angeletos and Lian (2018) and Farhi and Werning (2019)), and behavioral models (e.g. Arifovic, Schmitt-Grohé and Uribe (2018) and Goy, Hommes and Mavromatis (2020)). A common justification advanced by many of these studies is that there is ample evidence to support their modeling choices from laboratory experiments.

This paper seeks to unify key elements of these alternative approaches by marrying adaptive learning and level-k reasoning in a single heterogeneous expectations behavioral model. Adaptive/heuristic learning and heterogeneous expectations capture the well-documented behavior of laboratory subjects in Learning-to-Forecast Experiments (LtFE), e.g. Hommes, Sonnemans, Tuinstra and Van De Velden (2007), Hommes (2011), and Hommes (2013).[1] At the same time, level-k reasoning 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), Crawford, Costa-Gomes and Iriberri (2013), and Mauersberger and Nagel (2018).

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 may be history dependent. 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.[2]

We assume agents agree on the level-0 forecast, are free to choose between k-level forecasts, and we allow for heterogeneity in their choices. Over time the

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

[2]In the limit, the deductions comprise a key feature of the eductive learning framework of Guesnerie (1992) and Guesnerie (2002).

proportions of agents using different k-level forecasts evolve in response to the size of recent k-level forecast errors. The ability of agents to adjust their depth of reasoning over time means that the distribution of forecasts, and the average expectation, evolve via two distinct mechanisms: an adaptive process based on past data and a reflective process based on strategic considerations.

We begin our development of the unified model by deriving sharp analytic results within a simple static framework along the lines of Muth (1961). We then take this model to the laboratory. We test the core predictions of the unified model using a standard LtFE design. Finally, we demonstrate the usefulness of the unified model to address two macroeconomic issues of interest: the sacrifice ratio and the power of forward guidance.

Our theory is developed using the univariate model $y_t = \gamma + \beta \hat{E}_{t-1} y_t$, where $\hat{E}_{t-1} y_t$ is the average of individual forecasts. The setup can be thought of as either a cobweb market model or a repeated beauty contest/guess-the-average game, where β is the feedback from expectations to economic outcomes. Thus we consider both positive and negative expectational feedback cases.

We assume that the model structure is known to agents. Using this framework, we seek to answer three questions. First, under what conditions will agents with initially heterogeneous expectations converge to the unique RE equilibrium (REE) of the model? Second, if convergence to the REE is obtained, how quickly does it occur and what happens to individual forecast heterogeneity during the transition period? Finally, how do the heterogeneous agents respond to anticipated events in the form of announced structural change, and what effect does this have on subsequent convergence to the new REE?

The properties of this model under adaptive learning, eductive learning, and RE are well-known. A unique REE exists, and if $\beta < 1$ the REE is stable under a wide range of adaptive learning algorithms. However, coordination on the REE via eductive reasoning only obtains if $|\beta| < 1$.[3]

We show that our unified model inherits $|\beta| < 1$ as one key condition that delineates distinct types of behavior. When this condition is satisfied, starting with an arbitrary distribution of agents using differing level-k forecasts, convergence over time to the REE is obtained. Importantly, convergence can be much faster than predicted by adaptive learning, though it is not instantaneous. Faster convergence can arise because there are two channels through which realized y_t affects $\hat{E}_t y_{t+1}$: through the level-0 forecast and through the proportions of agents using different k-levels. Quick convergence can result in the coexistence of high-level

[3]The eductive approach is based on strong common knowledge assumptions of both the economic environment and the rationality of other agents. See Guesnerie (1992) for application of this approach to the cobweb model.

and low-level reasoners for extended periods: once the market has approximately converged to the REE, all level-k forecasts can provide almost the same prediction, greatly reducing the incentive to revise k upwards. Only in the limit, and in the absence of structural change, will low-level reasoners be driven out the economy.

We emphasize that agents sometimes revise downwards their reasoning level k. For example, consider a structural change after the market has approximately converged to the REE. In the presence of both high-level and low-level reasoners the endogenous response to structural change is muted relative to the RE prediction. Consequently, the highest level reasoners will make large forecast errors, and many of those high-level reasoners will find it optimal to reduce their reasoning level.

In the $\beta < -1$ case, the unified model makes other novel convergence predictions. Convergence to the REE is possible. Unstable dynamics are possible. Bounded cycles that are not centered at the REE are also possible. Which asymptotic pattern obtains depends on how quickly level-0 forecasts are updated, how quickly agents revise their depth of reasoning in response to forecast errors, and the initial distribution of level-k types.

The unified model is able to explain the well-documented phenomenon in LtFE that markets can converge much more quickly to the REE than predicted by adaptive learning alone. Our model also directly addresses the results of the Bao and Duffy (2016) experiment, which appeared to indicate a mixture of both adaptive learning and eductive reasoning. They note specifically that, when the condition for eductive stability – $|\beta| < 1$ in our framework – is not satisfied, market dynamics are distinctly different: both stable and unstable markets are observed, consistent with our unified model.

In our LtFE, we adapt the experimental design of Bao and Duffy (2016) to test the key predictions of the unified model. We place laboratory subjects into a computer-based market that nests the cobweb model. The market price depends on expectations, which are supplied by experimental participants who have full information of the market structure including the exact equations that govern supply and demand. Participants are paid based on the accuracy of their forecasts. In contrast to Bao and Duffy (2016) we consider positive as well as negative expectational feedback cases.

A key novel dimension that we add to the experiment is announced structural changes at irregular intervals. The announcements are akin to embedding a beauty contest game within a sequential market game, allowing us to see how participants incorporate new information into their forecasts. A major advantage of this approach is that the periods leading up to an announcement provide data for participants - and for us as the researchers – that clearly identify the level-0 beliefs from which level-k forecasts are derived. The unified model then provides sharp predictions for the distribution of forecasts observed in announcement periods, as well as for how people should revise their depth of reasoning in subsequent periods.

The unified model captures well the experimental data. 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 how we 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.

A shared history of market play therefore appears to coordinate subjects on a shared level-0 forecast, which either triggers level-k behavior or simply makes it easier to observe in the lab. This finding has implications for tests of level-k reasoning in settings when there is no clear level-0 forecast, e.g. the survey evidence in Coibion, Gorodnichenko, Kumar and Ryngaert (2021), which finds mixed evidence for level-k reasoning when settings vary.

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. This behavior 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 the key prediction of the unified model that revisions to depth of reasoning are not monotonic in the wake of announcements. We document, as the model predicts, that some high-level reasoners experience large forecast errors in announcement treatments because of the presence of low-level reasoners. This causes a fraction of the high-level reasoners to revise down their depth of reasoning. These downward revisions make the prevalence of low-level reasoning very persistent, providing support for macroeconomic models such as Farhi and Werning (2019) or Angeletos and Sastry (2021) that rely on low-levels of deductions to continually moderate general equilibrium effects.

Finally, we turn to the macroeconomic applications: the cost of disinflation and the power of forward guidance. First, we show that the unified model provides a mechanism for understanding the wide range of sacrifice ratios observed across countries and across disinflation episodes within countries. Second, using a simplified New Keynesian environment, we show that forward guidance policy is much less powerful under the unified dynamic than under RE. As in the laboratory experiment, announced policy induces some agents to revise down their depth of reasoning in response to forward guidance, lowering its overall impact.

Related Literature. We build off a large literature in theoretical macroeconomics on adaptive learning and rationally heterogeneous expectations models. We draw on the basic notions of expectation formation and stability, laid out in the adaptive learning (AL) literature developed in Bray and Savin (1986), Marcet and Sargent (1989), Evans (1989) and Evans and Honkapohja (2001), to ground the level-0 forecast assumptions and analysis. AL is a versatile technique that has been applied in a wide range of both nonexperimental and experimental settings. Importantly, AL is applicable whether or not agents fully know the structure of the economy. For a wide range of models it has been shown that agents acting as econometricians, and using least squares to update the coefficients of their forecast rules, can learn over time to have RE.[4] Under AL agents do not need to know or estimate structural parameters of the model, but can simply regress the variable(s) they need to forecast on an intercept and any relevant observed information variables, and update over time the forecast model coefficients.

Even when agents know the full economic structure, including structural parameter values, as in the present paper, they still will be uncertain of general equilibrium outcomes because they do not know the expectations of others. A large experimental literature has shown that subjects often use relatively small k-levels, and that there is heterogeneity across agents in their choices. AL is a natural anchor for forecasting equilibrium outcomes when the depth of reasoning of other agents is unknown.

To complete our model we draw on the behavioral heterogeneous expectations literature Brock and Hommes (1997), De Grauwe (2012) and Hommes (2013), which considers ex ante homogeneous agents selecting from a menu of forecast rules, resulting in ex post heterogeneity in a variety of macroeconomic settings. We differ from these treatments by considering menus with arbitrarily many k-level forecast types. Agents in our framework are free to scale up or down their depth of reasoning as they see fit, based on recent forecasting performance. Unlike this earlier literature or the calculation equilibrium approach of Evans and Ramey (1992) and Evans and Ramey (1998), we do not require calculation or forecasting costs. Agents choose not to jump immediately to the REE forecast because of their recognition that other agents are not (yet) making RE forecasts.

Our model shares elements with the Reflective Equilibrium notion proposed by García-Schmidt and Woodford (2019), which features heterogeneous level-k reasoners. However, their analysis takes place at a single point in time, and

[4] Indeed, as stressed, for example, by Sargent (2008), REE are most plausibly viewed as an emergent outcome from learned behavior. Hence, transitional learning dynamics, especially following structural change or a change in policy, can be important.

studies the implications of a finite degree of reflection for the current aggregate variables and for the average expected path of future aggregates variables. In contrast our unified approach specifies real-time dynamics for the time paths of level-k forecasts and the proportions of agents using each forecast level, as well as for the associated path of y_t.

The Cognitive Hierarchy (CH) approach of Camerer, Ho and Chong (2004) allows for a distribution of k-step types. The k-level reasoning used in our unified model differs from their approach in several crucial ways. In the CH framework the distribution of k-step types 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 these 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 come to know the distribution of lower-level types, yet at the same time not realize that there are other agents using equal or higher reasoning steps. These lacunae in the CH framework would appear even more salient in extensions to repeated or dynamic games.

In contrast to the CH approach, our unified model generates an evolving sequence of k-levels that relies on natural bounded rationality considerations. In a dynamic setting level-0 forecasts reflect observed prices to date. Given the level 0 forecast, standard level-k forecasts can easily be computed recursively. Agents do not need to observe, know or make assumptions about the distributions of level k forecasts used by other agents. Instead, agents revise their k-levels over time based on observed relative forecasting performance. This approach combines adaptive learning, computation of implied k-level forecasts, and revisions of agents' choices of k-level in light of actual forecasting performance.

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) 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 k-level forecasts are therefore not studied. 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 k-level 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. They observe heterogeneous forecasts with different depths of reasoning, consistent with our findings and with the unified model.

In Anufriev et al. (2022), subjects must forecast two variables whose realizations 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 is designed to test these predictions.

2 The Model

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

2.1 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 $\mathbb{N}$.[5] For $k \in \mathbb{N}$, the

[5] To be clear, $\mathbb{N}$ is assumed to include zero.

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* $\omega = \{\omega_0, \dots, \omega_M\}$, 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.[6] 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 3.

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

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

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).[7] 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 the 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.[8]

The most natural level-0 belief will depend on the model. For example, the

[6] It will sometimes be convenient to interpret a weight system as a sequence with only finitely many nonzero terms.

[7] It is also possible to extend the model to include additional types of heterogeneity. For example, agents could hold heterogeneous expectations over the level-0 forecast. Or, heterogeneity of the distribution of k-types could be taken into account by the agents at every level such as in the cognitive hierarchy model of Camerer et al. (2004). We view the level of sophistication and the degree of heterogeneity as an empirical question, which we study with an LtFE in Section 4 and 5.

[8] 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 $k \geq 1$.

level-0 belief may reflect a salient value, as in the guessing game model in Nagel (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).

2.2 Adaptive dynamics

We define *adaptive dynamics* as corresponding to adaptive learning with fixed level-k weights.[9] 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},

[9] 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 $i \geq 1$.[10] 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.

2.3 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.[11]

[10] 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 5 – 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.

[11] Note that the existence of the arg min is guaranteed by the fact that if $|\beta| < 1$ then

Next, assume the presence of a rate function $r : [0, \infty) \to [\delta, 1)$ with $\delta \geq 0$ satisfying $r' > 0$.[12] 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 $\omega_0 = \{\omega_{k0}\}_{k \in \mathbb{N}}$.

2.4 Unified dynamics

Unified dynamics joins adaptive dynamics and replicator dynamics. The level-0 forecasts are updated over time as in Section 2.2 and the weights evolve according to the replicator as in Section 2.3. The system under unified dynamics is given as

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

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.

$E^k_{t-1} y_t \to 0$ as $k \to \infty$, and if $|\beta| > 1$ then $E^k_{t-1} y_t \to \infty$ as $k \to \infty$.

[12] 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.

3 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 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.

3.1 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.[13]

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. Some partial results are available under adaptive dynamics, as discussed below.

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.

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

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

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

The earlier-defined set of all weight systems, Ω, is the disjoint union of these simplexes: $\Omega = \dot{\cup}_n \Delta^n$, where the dot over the union symbol emphasizes that, 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.[14]

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 are also able to draw conclusions when $\beta < -1$. The surjectivity of ψ results from the expanding magnitudes and oscillating signs of the β^n. The 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 $\beta < -1$.*

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

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.

3.2 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 the *aggregate parity*, which can be viewed as an aggregate measure of optimism and pessimism.[15]

3.3 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 $k = 0, 1, 2, 3$.

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

Figure 1: Simulated dynamics with positive feedback

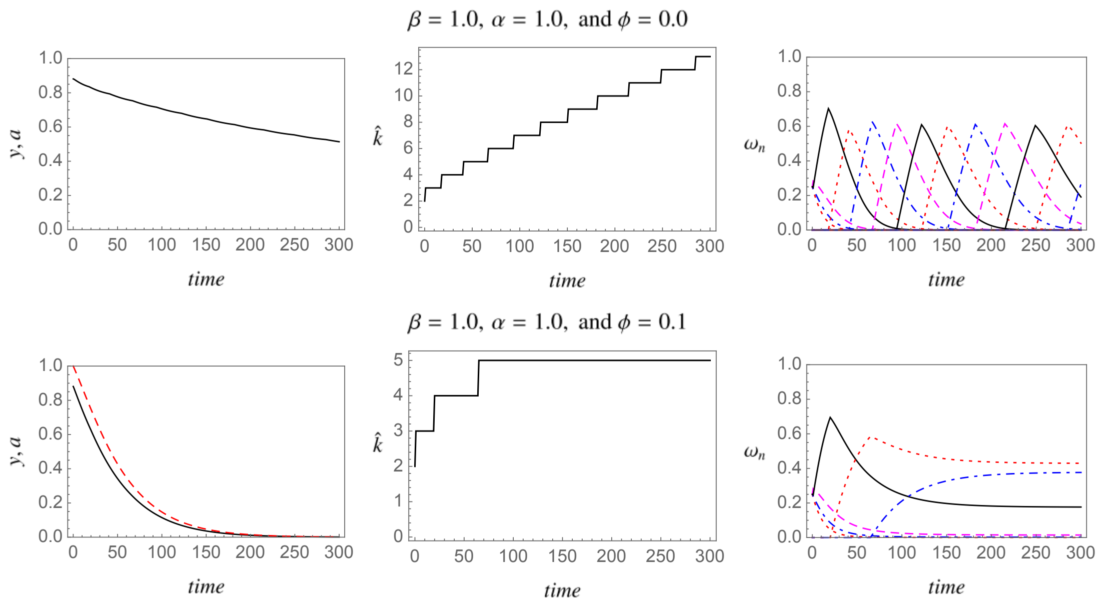

Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). In the left panels the *solid black* curves denote y and in the bottom left panel the *dashed red curve* identifies a. In the right panels $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights 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 1, 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 2 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, $\hat{k}$

Figure 2: Simulated dynamics with negative feedback.

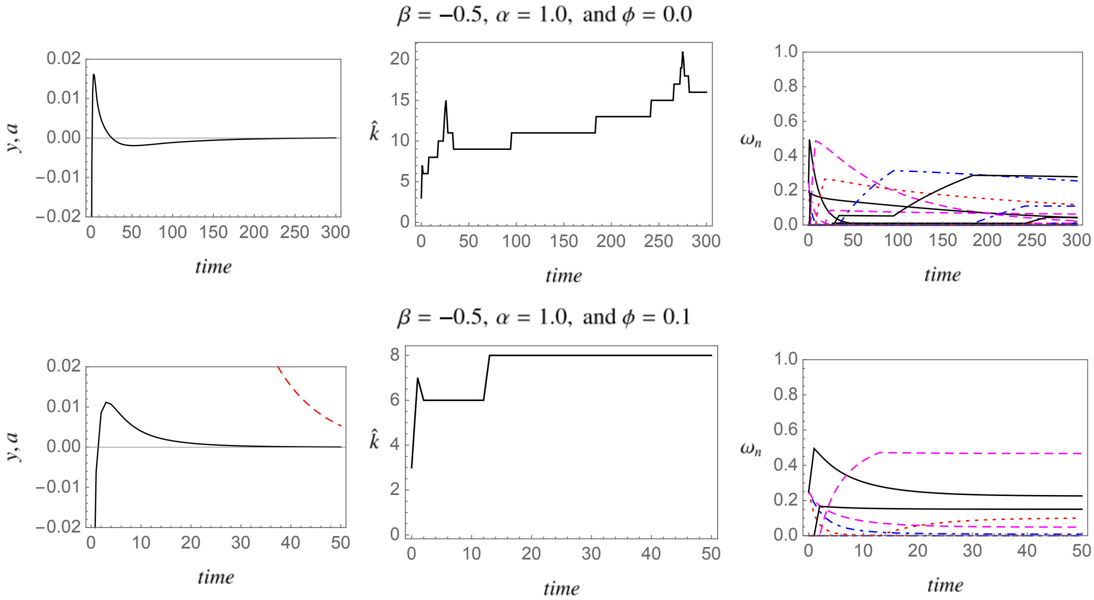

Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *red dotted, blue dash-dot, dashed magenta* and *solid black*, respectively.

rises sharply to drive down (in magnitude) the optimal forecast $\beta^{\hat{k}}$.

By Theorem 2, $\hat{k} \to \infty$. However, unlike the positive feedback case, here this convergence is not monotone. Figure 2 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 3: Simulated dynamics with large negative feedback.

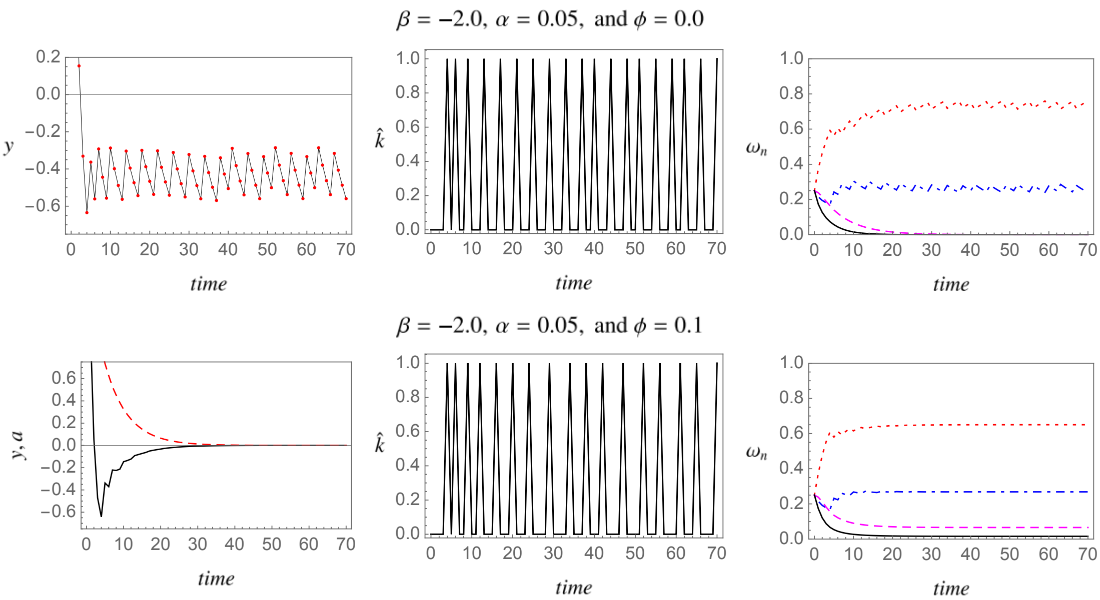

Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *red dotted, blue dash-dot, dashed magenta* and *solid black*, respectively.

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 Guesnerie (1992): 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 3 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.[16] The bottom row of Figure 3 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.

3.4 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, $\bar{y} > 0$. In the lab experiments we use a market model with free disposal, which precludes negative prices, and it is therefore useful in the current section to include non-negativity constraints on y and $E^k_{t-1}y_t$.

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.[17] 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.

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

[17] 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 6 examines the unified dynamic to models in which y_t depends on $E_t y_{t+1}$.

Figure 4: Unified dynamics with announced structural change in period 20 and 45.

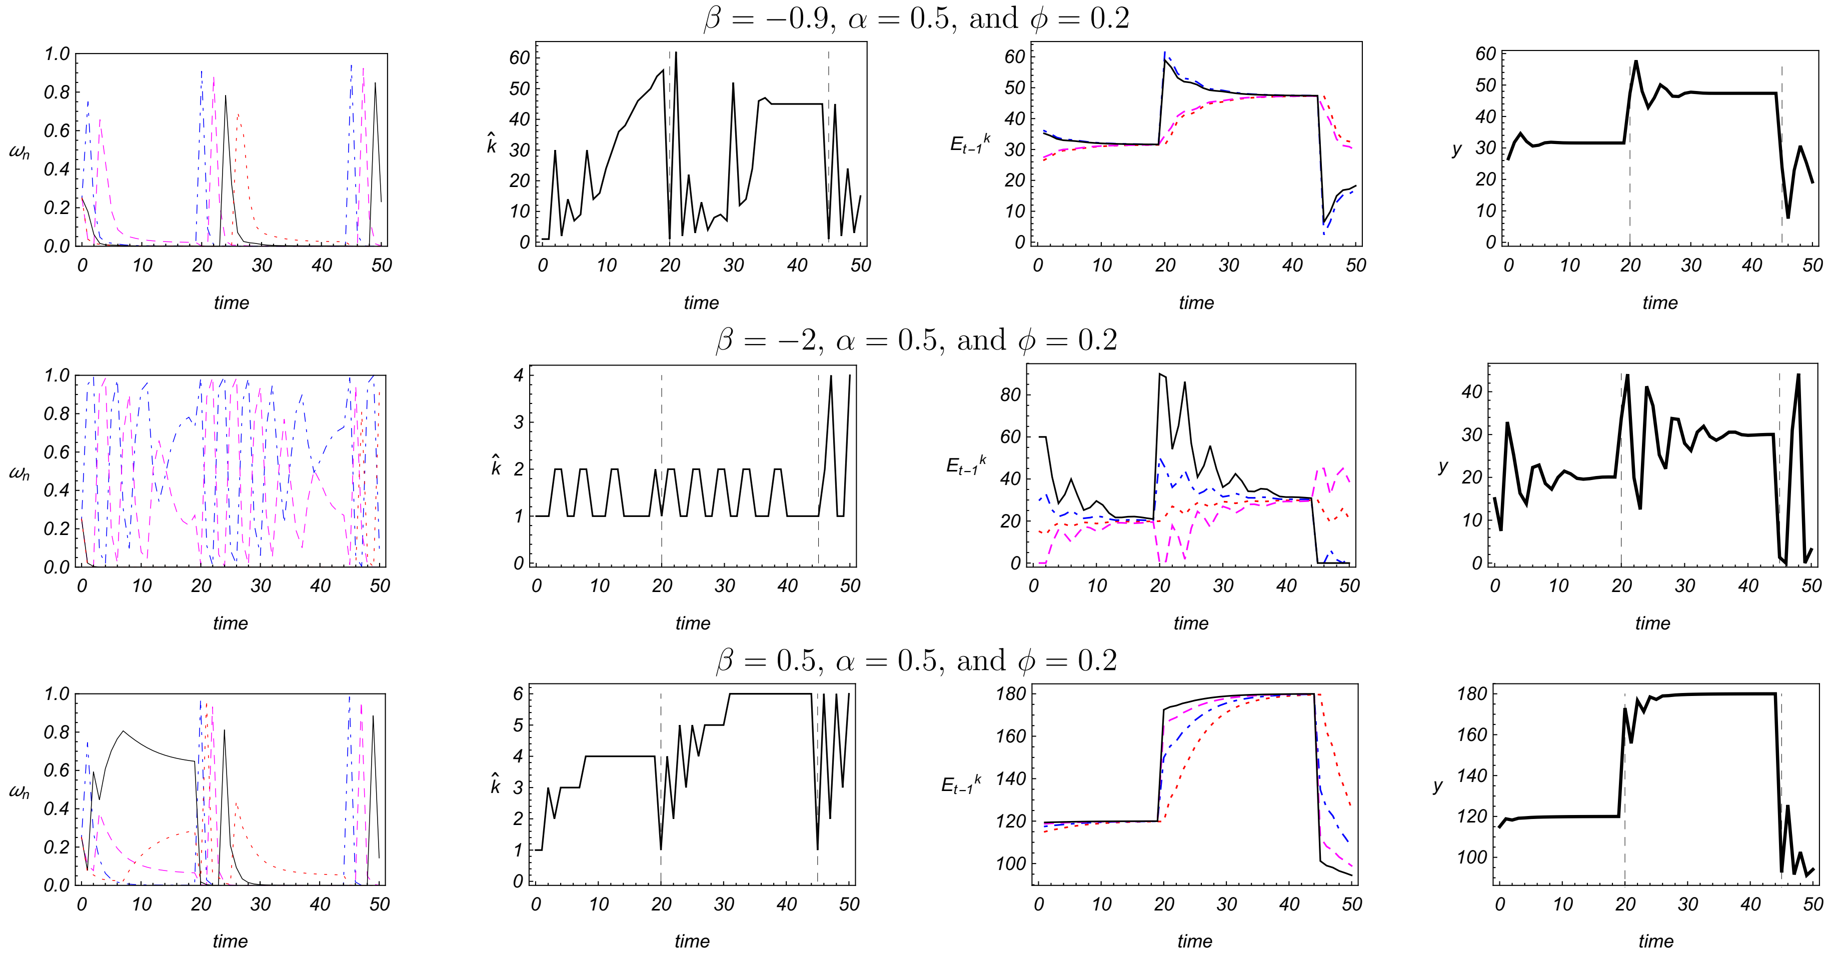

Notes: Simulation of unified dynamics with announced changes to the intercept and a known non-negativity constraint. $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *red dotted, blue dash-dot, dashed magenta* and *solid black*, respectively.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

Figure 4 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 additional 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 4 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 4. However, the mass of high-level reasoners generally increases over time with repeated announcements.

The middle row of plots in Figure 4 shows a simulation for $\beta = -2$. Here the choice of parameters matters greatly for the outcome, and we consider a

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 $t = 45, ..., 50$.

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.

4 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 3.4 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. The market price is determined 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 experimental settings. The shock sequence is the same in all markets and across all 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.

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 3.3 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.[18]

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 4 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 level of reasoning to the current optimal predictor.*

The unified model posits that not all agents revise in every period, but those that do revise to the optimal predictor based on the last period's price. This implies that revisions to the depth of reasoning may be non-monotonic in some instances. In particular, following announcements, we expect revisions to the depth of reasoning for those agents who experience large forecast errors as shown in Figure 4. Additionally, our theoretical and numerical results suggest the following hypothesis.

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.*

[18] 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).

We prove that convergence under the replicator dynamic is obtained when $|\beta| < 1$ through increasing k-levels of reasoning. When $\beta < -1$, low k-level strategies must be maintained or else market prices do not converge.

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.

4.1 Experiment description

The experiment used a computer based market programmed in oTree.[19]. 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. Laboratory participants were told that they are acting as expert advisers to firms that produce widgets. Participants 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.

We checked for comprehension of the market environment with a version of the following question in the tutorial:

> Consider the case where $A = 60, B = 2, D = 1$ and $noise = 0$. If we substitute these numbers into this equation
>
> $$p = \frac{A}{B} - \frac{D}{B} \times \text{average price forecast} + noise,$$

[19] See Chen, Schonger and Wickens (2016) for documentation.

Figure 5: Screenshot of experimental market GUI

we get that price (p) is

$$p = 60 - \frac{1}{2} \times \text{average price forecast.}$$

What is the market price (p), if the average expected price is equal to 38?

Participants were not able to continue with the experiment until the question was answered correctly.[20] A worked version of this problem with different numbers was also provided on the printed instruction sheet. The question was designed to verify that each participant knew how to use the equations without teaching the person to solve for the REE. The tutorial and printed instructions are available in the online Appendix A7.

Figure 5 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 the past prices, predictions, forecast errors, and the forecast's earnings is provided on the left-hand-side of the screen.

[20]Four out 372 participants were not able to solve the question on their own and asked for help from the lab manager. In this case, they were directed to look at the example on the instructions, which clarified the problem in all cases.

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. Negative quantities receive zero cents. The function is presented and explained to participants as part of the tutorial and is the same across treatments. Forecasts must be within 4 units of the actual price to earn money for a forecast. We chose this specification to give participants a high incentive to be precise in their predictions when confronted with announcements. Previous studies have employed point systems that compensate more generously for poor forecasts. For example, in Bao and Duffy (2016) participants needed to be within 7 units to earn points, which ranged from zero to 1300.

In addition to performance pay, subjects received a \$5 show-up fee. In the T2 treatments, subjects also received an additional \$5 of guaranteed compensation to offset the lower earnings that we expected (and which did occur) in these treatments due the difficulty in coordinating.[21] The difference in guaranteed pay and the treatment settings were not disclosed to the subjects in advance.

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.[22]

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

5 Experimental Results

Table 2 reports summary statistics for the experiment. In total, 372 individuals participated in 62 experimental markets. All T1 and T2 treatments were conducted in May and June of 2018 at the UNSW Sydney BizLab. Two sessions for each treatment were scheduled with the aim of testing at most five markets in each session. Participant no-shows account for the different number of markets

[21] Ethics requirements placed on the study mandated that participant payments were on average \$15 AUD per hour.

[22] A minimum price of 0 and a maximum price of 500 was enforced as well. These bounds were not advertised to participants but if chosen a pop-up box would appear informing them of the bound.

Table 2: Summary statistics

Treatments	Markets	Participants	Treatment Values		Payments		Time Use (min)	
	(62)	(372)	Feedback	Annoucements	Total Pay	Pay Efficiency	Tutorial	Total
T1 x A1	6	36	-0.9	1	\$20.31	81%	9.5	64.9
T1 x A2	7	42	-0.9	1	\$18.68	75%	7.4	71.0
T1 x A3	7	42	-0.9	2	\$17.76	71%	8.2	75.3
T2 x A1	7	42	-2	1	\$14.52	58%	8.8	66.7
T2 x A2	7	42	-2	1	\$13.30	53%	8.2	84.9
T2 x A3	8	48	-2	2	\$11.17	45%	7.6	80.4
T3 x A2	9	54	0.5	1	\$17.62	70%	8.2	57.2
T3 x A3	11	66	0.5	2	\$18.18	73%	8.2	61.7

Notes: Pay efficiency is the total possible pay for accurate forecasts divided by the maximum pay of \$25 per session, which does not include show-up payments or top-ups.

across the treatments.[23] All T3 treatments were conducted in March of 2019 at the University of Sydney's Experimental Lab. Eight sessions were held with the aim of testing at most four markets in each session.[24] Again, no-shows account for the different number of experimental markets across treatments.

Figure 6 provides a visual overview of the experimental results from the T1×A3, T2×A3, and T3×A3 treatments. These three treatment illustrate the most novel features of our experimental results and provide a qualitative comparison to simulated unified model in Figure 4. The first column of figures shows the proportion of laboratory participants that we identify as level-0, 1, 2, and 3 in each period. We provide the exact details of this classification in Section 5.2. The second column shows the distribution of participant's forecast types that we observe in the first announcement period (round 20) compared to the second announcement period (round 45). We discuss these results in Section 5.3. The third column shows the median forecast of participants that we identify as level-0, 1, 2, or 3 shown in the first column of plots. The final column shows the average market prices from the experimental markets and the individual forecasts.

5.1 Convergence Results

The last column of Figure 6 illustrates general convergence properties found across T1 - T3. 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.

[23]Subjects were recruited using ORSEE (see Greiner, 2015).

[24]The different session sizes at the University of Sydney reflect lab capacity constraints due to equipment issues and subject recruitment limitations. Christopher Gibbs left UNSW for Sydney in July of 2018, causing a the delay between experiments.

Figure 6: Comparing the unified model to experimental data

T1 × A3 ($\beta = -0.9$)

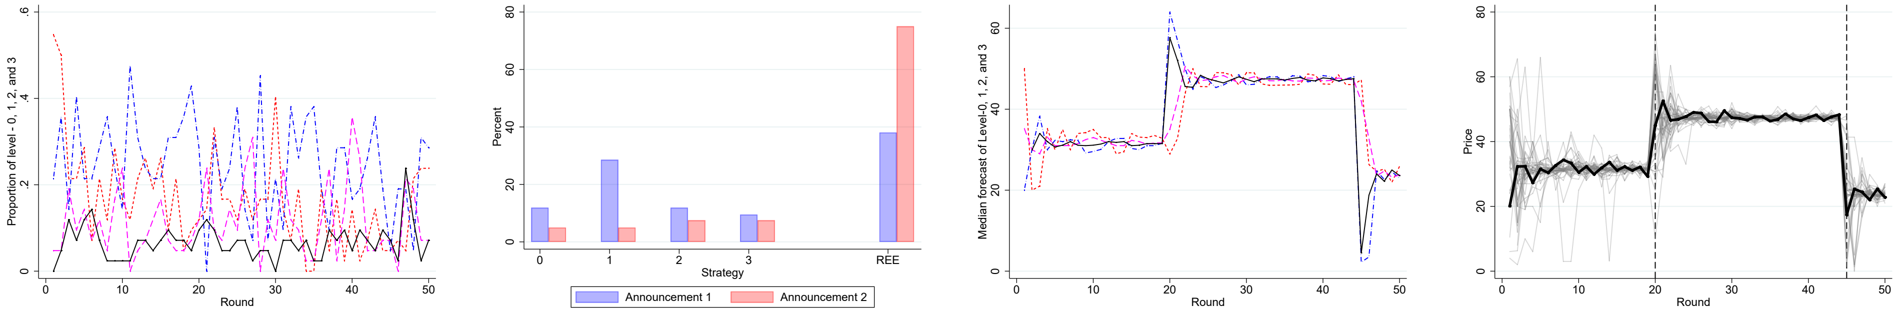

T2 × A3 ($\beta = -2$)

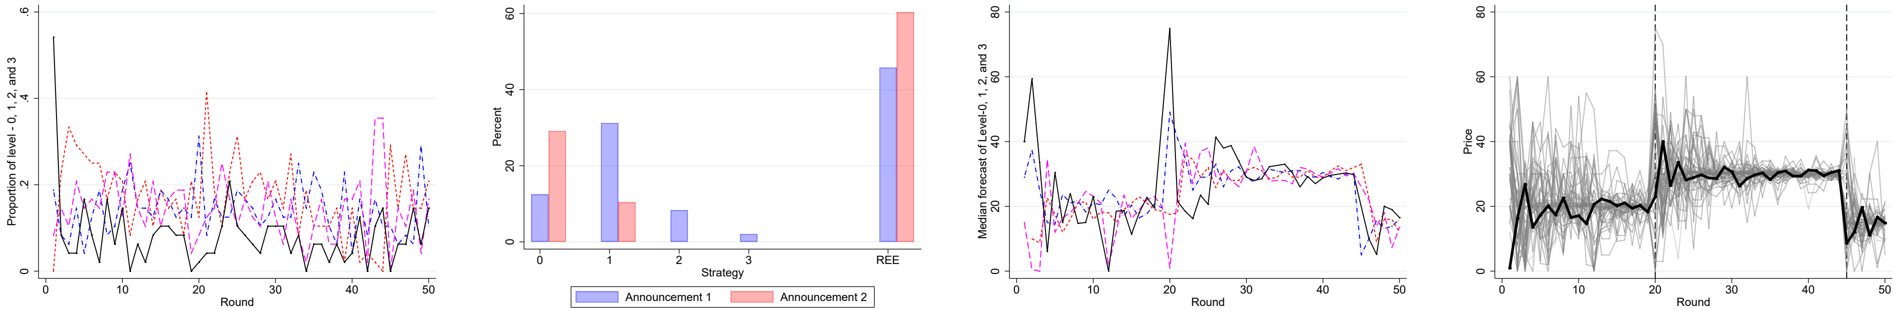

T3 × A3 ($\beta = 0.5$)

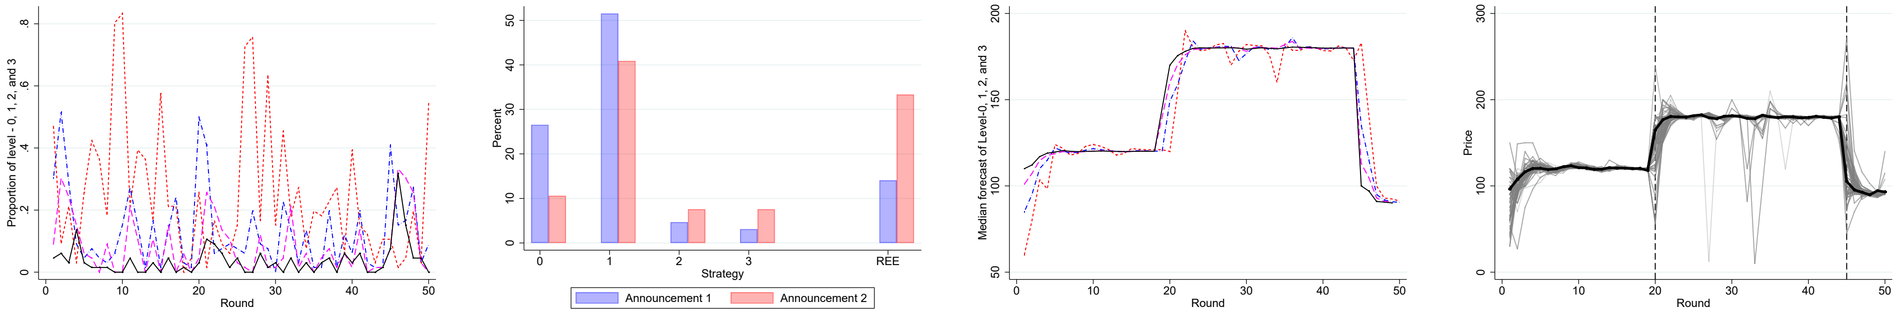

Notes: Survey participants' forecasts are classified as Level-0, 1, 2, 3, or consistent with the REE forecast by comparing to the model implied forecasts. The time path of observed ω_n for $n = 0, 1, 2, 3$ are distinguished by plot-style: *red dotted, blue dash-dot, magenta dash* and *black solid*, respectively. The corresponding median forecasts, $E_{t-1}y_t^k$, of the participants use the same styling. The second column shows the distribution of the types of forecasts observed in the announcement periods. The final column shows average market prices observed (solid black) laid over all individual forecasts.

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. First, because of the random noise component of price, we define a round to be converged when the price is within plus or minus three of the steady state price. Based on this cutoff, we simply count the number of rounds in a given interval in which price satisfies this criterion. Columns 2 - 4 of Table 3 shows the count data for the three feedback treatments, where we look at various 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 at least 85% of the rounds satisfy the above convergence criterion. 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 second metric we use to assess convergence is the mean difference in the market price from steady state over the same intervals used for the first metric. Columns 5 - 7 of Table 3 show the mean difference and the t-statistics for a test of the null hypothesis that the mean difference is less-than-or-equal to 3. Bolded values indicate a one-sided rejection of the null hypothesis with a p-value smaller than 0.15. By this metric, convergence is achieved in the T1 treatments within 5 rounds of an announcements and maintained through all other intervals. Convergence is achieved for the T3 treatments in rounds 6-10, but within five rounds after the first announcement. A t-test of the difference in this measure for rounds 2 through 19 versus 21 through 38 confirms that market prices are closer to steady state after the first announcement (bottom row of Table 3) than at start of the experiment, which indicates faster convergence after the announcement. The T2 treatments again show a different pattern. With this metric we only find marginal convergence for rounds 11 - 19 and 31 - 38 in treatments with an announcement. But we do find that prices are on average closer to steady state following the announcement.

The final metric we use to assess convergence is the average earnings by participants per round over the same intervals previously studied. The maximum earnings in a round is $0.50 and forecasts must be within plus or minus four of the actual price to earn money. Therefore, high average earning indicates that all market participants are making accurate forecasts. The last three columns of

Table 3: Convergence of price to REE in experimental markets

Rounds	Ratio of Market Rounds Converged (Converged/Total)			Mean $\|p_t - \bar{p}\| = \mu$ $H_0: \mu \leq 3$ $H_a: \mu > 3$			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.

Table 3 show the mean earnings and the t-statistics for a test of the null hypothesis that average earning are greater-than-or-equal-to $0.40. This is our strictest measure of convergence. By this measure, we only observe convergence in rounds 11 - 19 and 31 - 38 for T1 and T3 treatments. We never observe convergence for the T2 treatments.

5.2 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 can be viewed as 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.[25] 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 subject's 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 levels of deductions as well.

Table 4 summarizes the proportion of individuals we classify as each type in each of the announcement rounds on the left side using the ± 3 cutoff. The data from all treatments is pooled. The ranges in brackets below the classification percentages show the proportion of forecasts that we would classify as each type if we used a ± 1.5 cutoff or a ± 4.5 cutoff. Overall, 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 4 provides a logical check on our classifications. It is natural to think that higher levels of deduction require greater cognitive resources. Therefore, a person who makes a level-0 forecast may not spend as much time formulating a forecast as someone who makes a level-3 forecast. Therefore, if our classifications are actually identifying people who are making level-k deductions, then we should find some correspondence to the time spent deliberating on each decision and the depth of reasoning that we identify. To investigate this

[25] The results are robust to reasonable changes in the definition of level-0 forecast. In the online Appendix available on the authors' website, we reproduce all of our results under the level-0 assumption of the average of the previous four prices for comparison purposes. We also explore one market in detail in the online Appendix, which illustrates further how the classification works in practice.

Table 4: 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.

hypothesis, we estimate the following regression model:

$$d_{i,r} = \alpha_i + \omega_r + \sum_{k=0}^{3} \beta_k I(k=1)_{i,r} + \sum_{j=0}^{3} \gamma_k I(k=1)_{i,r} \times I(r=Ann)_r + \epsilon_{i,r} \quad (7)$$

where $d_{i,r}$ is the time spent deliberating for person i in round r, α_i is an individual fixed effect that controls for fixed characteristics such as treatment and unobserved fixed individual or market idiosyncrasies, ω_r is a round fixed effect to control for the fact that generally less time is spent deliberating in later rounds, $I(k=1)_{i,r}$ is an indicator that takes a one if we classify a person as choosing a level-k forecast in round r, and $I(r=Ann)_r$ is indicator variable for announcement rounds. We cluster our standard errors at the market level. The coefficients β_k and γ_k identify the difference in time spent for those identified as level-0, 1, 2, and 3 overall and in the announcement periods, respectively, relative to those we identify as choosing the REE forecast, or who 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 7 shows histograms of the individual forecasts in round one and each subsequent announcement round for each feedback treatment. The gray bars show the model implied level-k forecasts with a plus or minus three band. The T2 round 20/50 predictions provides the most clear level-k deductions because the significant negative feedback ($\beta = -2$) in the market makes each level-k prediction very distinct.[26] Overall, it is clear that there is not strong evidence for level-k reasoning in the first period. In fact, a significant fraction of the total number of people that we classify as level-k is due solely to our ex post choice for the level-0 forecast. Many laboratory subjects do not appear to understand the structure of the game in the first period. Most appear to select whole numbers without much strategic thought. However, that changes once participants have played multiple rounds and an announcement occurs. For these announcements, Figure 7 and Table 4 show a majority of participants playing level-k or the high level-k/REE forecasts. The exit surveys also provide support for this interpretation with on average participants claiming that the equations and a forecast of average expectations were more important to calculating their own predictions than they believed they were to other participants' calculations. Past prices were thought to be more important to others' forecasts than their own forecasts, indicating a beliefs that others behaved adaptively, consistent with our level-0 assumption. In the interest of space, the survey results are discussed in the online Appendix.

Finally, we return to Figure 6 that summarizes the overall dynamics we observe in the data. Here we do not make use of cutoffs for the classifications and instead use all of the data to summarize overall observed behavior. To do so, each individual forecast is classified as level-0, 1, 2, or 3 based on the whichever forecast it is closest to measured by absolute error.[27] The first column of the figure shows the proportions we identify as level-0, 1, 2, and 3 over time. The third column shows the median forecast from those we identify as each type. Even without narrowing our classification of type with cutoffs, the unified model provides a good prediction of median individual behavior observed among experimental

[26] In the online Appendix, we provide more evidence that subjects make oscillating deductions as the market converges, consistent with level-k reasoning.

[27] Although we do not impose a cutoff when classifying forecasts types in this exercise, we choose not to classify 35 out of the 18,367 forecast observations that are clear outliers. For example, these include cases where a market had been converged for many periods at a price of 30 and a participant entered a one-off forecast of 300. Many of these forecasts, without cut offs, would be classified as REE or level-0, which are the nearest predicted forecasts, and which is clearly not in keeping with the goals of classifying forecast types. We include a detailed discussion of outliers in the online Appendix.

participants, especially in announcement periods.

5.3 Revisions to the depth of reasoning

Revisions to the depth of reasoning via the replicator employs three key assumptions. First, it assumes that not every agent will update their forecast in every period. Second, the agents who do should on average experience larger forecast errors in the most recent period. And finally, a person's choice of a new strategy should be based on a counterfactual exercise, where alternative level-k deductions are evaluated on the most recent outcome, and the best strategy from this reflective process is selected.

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. They generate large forecast errors for many participants, and they provide distinct counterfactual level-k predictions, which we can use to identify subsequent revisions to the depth of reasoning in the experimental data. Specifically, comparing individual outcomes and predictions in the announcement rounds to the round following the announcement, we can assess who has revised their depth of reasoning, how the revision compares to the best level-k forecast one could have chosen in the announcement period, and whether those who changed strategy experienced larger forecast errors. To maximize the data and to not exclude those who decided to switch from a non-classified strategy to a level-k strategy, we do not impose a cutoff when classifying a person's forecast as level-0, 1, 2, 3, or the REE for this analysis. Classifications are made based on whichever level-k strategy the submitted forecast is closest to in mean squared error.

Figure 7: Laboratory subjects' forecasts in announcement rounds

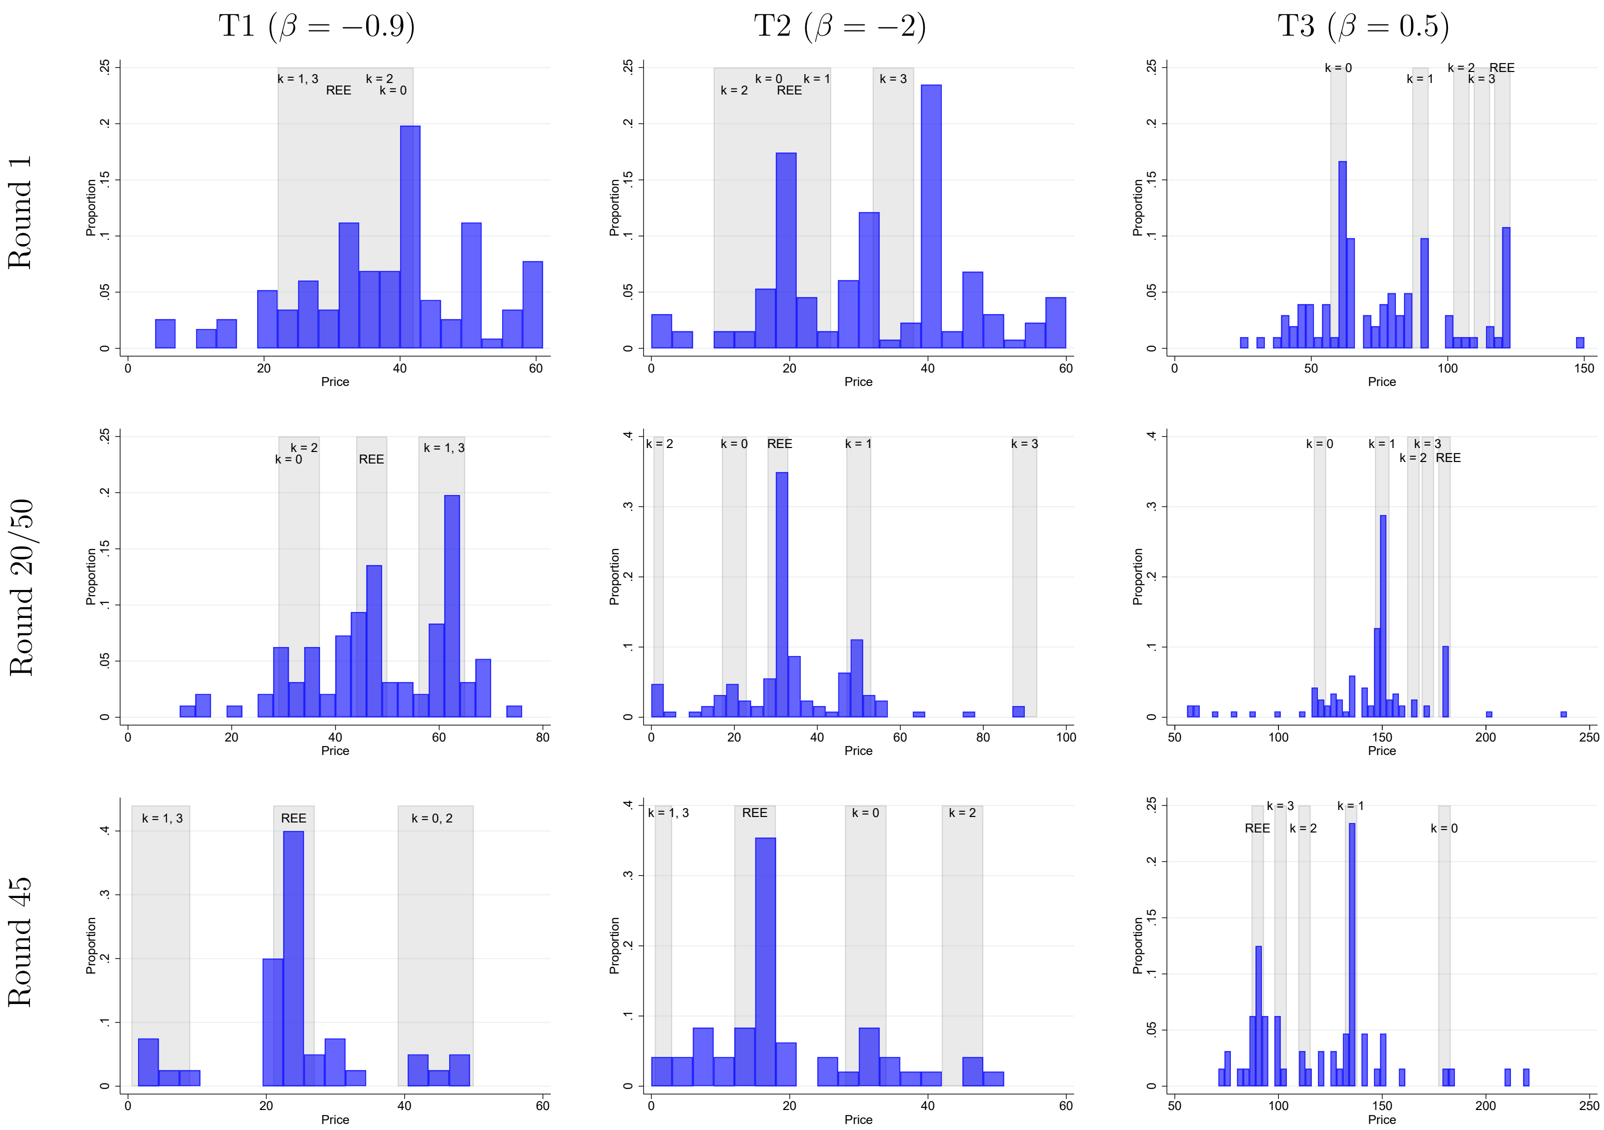

Notes: Histograms of the subject's forecasts in response to an announced structural change. The shaded regions correspond to our classifications of level-0, 1, 2, 3, and the REE forecasts reported in Table 4, which is ±3 of the model implied Level-k forecast. The width of each bin for the experimental 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.

Table 5 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, are classified as selecting the best counterfactual strategy from the previous period, which was often a lower level of reasoning than the one played in the announcement round as predicted by the unified model. 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 indeed do not update their strategy following the announcement period. Second, the subjects we do classify as changing strategy on average had experienced larger forecast errors and subsequently spent more time deliberating compared to those who did not change their strategy. Only in the T3 x A2/A3 treatment do we not find full congruence to the predicted pattern. In this treatment, changers make larger forecast errors, but spend less time deliberating. However, the difference in deliberation time is not statistically significant. Finally, of the subjects who we observe changing strategies, a significant proportion are classified as changing to the strategy that would have been the best level-k strategy from the previous period. The proportions we document here are significantly larger than what one would expect to occur by chance in all cases except for the T2×A3 treatment.

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. Figure 6 shows the distribution of forecasts for levels-0 to 3 and REE in the second column of plots. 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.

Table 5: 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.

5.4 Quantitative Evaluation

The previous sections focused on whether the experiments confirm the individual-level behavior predicted by the unified model. In this section we explore the fit of the unified model to the aggregate price data relative to alternatives models. In other words, from a macroeconomic perspective, how much better, if at all, does the unified model predict market price dynamics?

To address this question, we fit alternative models to the experimental data using a simple statistical approach. We consider four competitors to the unified model: REE, a fixed level-k model, a replicator-only model, and an adaptive learning model. For each model (except REE), we conduct a grid search over the relevant forecast parameters to minimize the squared error between the simulated data and the experimental data.

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 but the proportion of agents using different level-k types is 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 model ignores the underlying heterogeneity and uses (3) as the forecasting rule for all agents.

Table 6 shows the average mean squared error (MSE) for the T1×A3, T2×A3,

Table 6: 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 $\phi \in [0, 1]$.

and T3×A3 markets. In the interest of space, the individual market outcomes are reported in the online Appendix. Wilcox ranked signed tests of the paired market MSEs reveal that the unified model statistically improves upon the adaptive learning model at the 10% level across the different feedback treatments and it beats the RE forecast at the 5% level for the T1 and T3 treatments. However, it fails to reject the null hypothesis for equality for the T2 treatments. We discuss this last result below.

The unified model also outperforms the fixed level-k model and the replicator-only model across all three treatments at at least the 5% level. This result, along with the improvement over adaptive learning alone, provides evidence that all three elements – adaptive learning, level-k reasoning, and the replicator – are required to explain the aggregate data.

5.5 Discussion

The experimental evidence 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 $|\beta| < 1$. In addition, the speed of convergence measured in multiple ways appears to increase following an announcement treatment (see Table 3). Increases in convergence speed in treatments T1 & T3 are also supported by the increase in the depths of reasoning we observe among subjects when there are multiple announcements: see Figure 6.

We find strong support for Hypothesis 2 (level-k reasoning). We observe level-k deductions taking place in each of the announcement treatments with clear bunching around the k-level predictions in the histograms shown in Figure 7. Comparing the individual forecasts to the model implied forecasts in announcement rounds, we classify between 50% and 70% of subjects, depending on the chosen cutoff, as Level-0, 1, 2, 3, or REE. Our classifications also coincide well with the deliberation times we observe among participants, with level-0 participants spending 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 5). As predicted by our theory, many of those changes correspond to decreases in the subject's k-level depth of reasoning.

Finally, we find mixed evidence for Hypothesis 4 (level-k dynamics). We observe revisions over time in depth of reasoning for the T1 and T3 treatments. There were also more high level-k forecasts played for second announcements compared to first announcements, along with quicker convergence (see Figure 6 and Table 3). In addition, we do 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. The reduction in the depth of reasoning in favor of level-0 forecasts observed here is consistent with hypothesis 4. However, the increase in the fraction of people who choose the REE forecast is at odds with the unified model. This finding also explains why the RE forecast fits the aggregate price data fairly well in the quantitative evaluation of competing models discussed in the the previous sub-section.

We speculate that the high proportion of REE forecasts observed in the T2×A3 treatment's second announcement round may be due to the fact that in the experiment 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.

6 Applications to macroeconomics

While our theoretical and experimental focus employed a cobweb environment, with either positive or negative expectational feedback, unified learning dynamics can be implemented in a wide variety of macroeconomic models in which RE is normally treated as the benchmark assumption. To illustrate the unified model's potential, in this section we consider policy implications in two stylized settings: in a new-classical environment, we compare the sacrifice ratio under unified learning with that implied by RE; and in a new-Keynesian model we consider the efficacy of forward guidance under unified learning dynamics.

6.1 The sacrifice ratio

A major issue for monetary policy is controlling inflation. In the US, for the period since 1960, unacceptably high inflation rates arose in the late 1960s, 1973-74, 1979-1984 and 1990-92. Most recently, high inflation returned beginning around May 2021, with inflation significantly above the Fed's 2% target. The primary policy tool for reducing inflation is tighter monetary policy: reducing the money supply growth rate and increasing the policy interest rate. Historically, reducing inflation has usually required an output cost; in models with an expectations-augmented Phillips curve, this cost is measured by the cumulative reduction of GDP relative to trend over the period of disinflation. The "sacrifice ratio" is typically specified as the total amount of output lost, per percentage point reduction in inflation, associated with an orchestrated reduction in the inflation rate. Ball (1994) found that, empirically, the average sacrifice ratio has varied considerably across countries (ranging from 1 to almost 3). Gibbs and Kulish (2017), using more recent data, show that the ranges of sacrifice ratios across countries, and within countries for different inflation episodes, are even larger.

The sacrifice ratio depends crucially on how inflation expectations are formed. In principle, under RE, the sacrifice ratio can be zero in many standard monetary models, provided the policy-maker follows a credible policy. This view was forcefully advanced in a pair of papers in 1981 and 1982 by Thomas Sargent, "Stopping moderate inflations: the methods of Poincaré and Thatcher" and "The ends of four big inflations" (see Sargent 2013), in which he argued that credible disinflations need not have large output costs. In understanding this point it is important to distinguish between two components of central bank credibility: confidence that the central bank will follow through on its announced policy of monetary tightening; and confidence that this policy will achieve an immediate disinflation. The latter component is based on the RE hypothesis, and requires not just individual rationality and confidence that the central bank will follow through on its policy, but also confidence that other households and firms will believe that the policy will be fully successful.

In the face of the high inflation rates since Spring 2021, monetary economists and policymakers have revisited the issue of the output costs of disinflation. Bullard (2022a,b) argues, building on the experiences of monetary tightenings in 1983 and 1994-5, that if monetary policy reacts to above-target inflation by tightening sufficiently quickly and strongly, then inflation expectations can be kept under control, recession can be avoided, and the output costs of disinflation will be relatively low. This dynamic is evident in the results of Gibbs and Kulish (2017) where, for example, between 1982 and 1998, Australia experienced three sequential disinflation episodes with drops in inflation of more than 3%, with declining sacrifice ratios of 6.8, 2.5, and 0.5, respectively; Germany experienced two such episodes between 1981 and 1996 with declining sacrifice ratios of 6.4 and 3.6, respectively; or Japan which experienced two disinflation episodes between 1980 and 1995 with a greater than 3% change with declining sacrifice ratios of 5.4 and 4.9, respectively.

We consider the sacrifice ratio within the context of our unified model of learning to forecast. Within the context of a classical inflation model that is nested by our framework, we show that the unified learning approach can provide a natural explanation for a wide range in sacrifice ratios. Furthermore, our framework speaks to how, following an earlier successful, if costly, disinflation policy, the experience of the earlier policy change can lead average expectations to respond more closely in line with RE in subsequent policy interventions.

Our results can be obtained using a simple textbook New Classical model:

$$\text{AS: } q_t = \varphi(p_t - p_t^e), \;\; \text{AD: } q_t = m_t - p_t, \;\; \text{PR: } m_t = p_{t-1} + g.$$

Here q_t is log of the ratio of real GDP to potential real GDP. The first equation is the standard Lucas aggregate supply (AS) curve, where p_t is log price level and p_t^e is expected p_t.[28] The second equation is the quantity theory version of aggregate demand (AD), and the third equation is the policy rule (PR). An alternative interpretation of AD, stressed in Woodford (2003), Ch. 3, is that monetary policy is specified in terms of a target path of the log of nominal GDP m_t which is achieved by changing the policy interest rate as needed. The policy rule is an endogenous nominal GDP growth rule, in which, given last period's price level, nominal income is expanded at a rate such that inflation $\pi_t = p_t - p_{t-1}$ would equal g if the output gap q_t were zero.

Combining equations gives the reduced-form model $\pi_t = \beta\pi_t^e + (1-\beta)g$, where $\beta = \frac{\varphi}{1+\varphi}$. The perfect foresight REE is $\pi_t = g$, with (log) real GDP normalized to zero.

Rather than add to our model exogenous inflation shocks that raise inflation above a fixed policy target, it is more convenient to illustrate the mechanisms of our unified model of learning by considering an experiment in which the central bank tightens policy to reduce inflation from a high level to a lower target level, and then, subsequent to reaching this lower steady state, it announces a policy to

[28] The AS curve can alternatively be derived under the assumption that there is a fraction of sticky price firms that each period set prices in advance. See Woodford (2003), Ch. 3.

further reduce the inflation target. Specifically we consider the following policy experiment:[29] with money growth and inflation at 10%, the economy is in a steady state. In period $t = 10$ the growth rate is reduced by 25% (i.e. to 7.5%), and then at $t = 40$ it is reduced another 25% (i.e. to 5.6%). Under RE, log GDP sits at zero and the inflation rate instantly adjusts. Figure 8 presents the outcome under the unified model of learning.

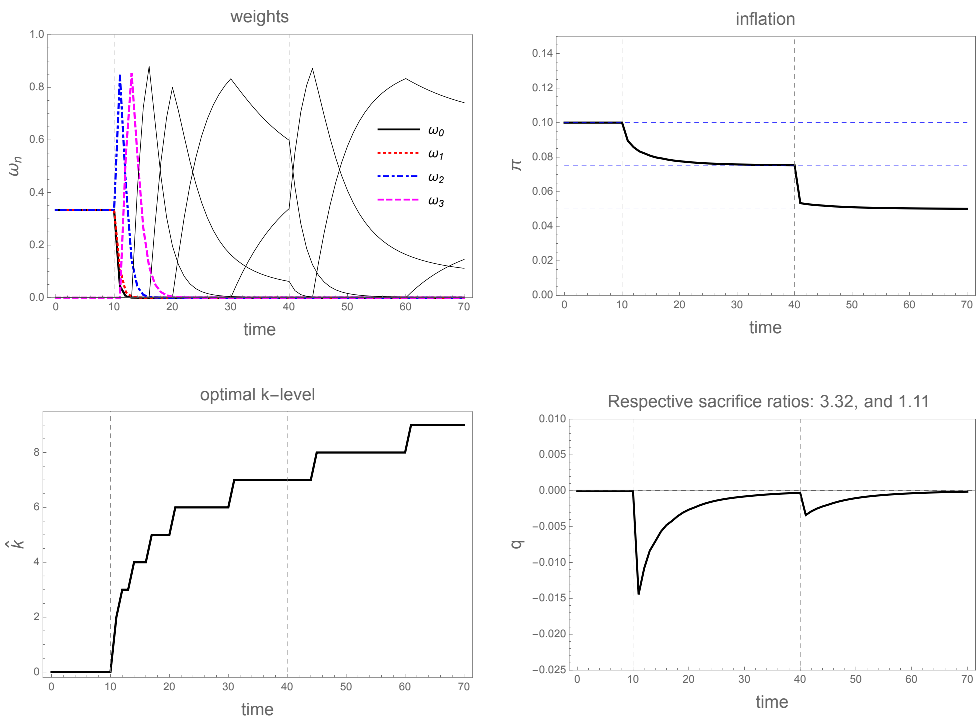

Figure 8: Sacrifice ratio

These figures illustrate the mechanisms that lead to enhanced policy "credibility." During the first disinflation, level-0 expectations gradually adapt to a lower inflation level, but also, importantly, the proportion of agents using higher level expectations grows over time because of their increasing forecasting success. As a result, when the second policy disinflation is implemented at $t = 40$, a larger proportion of agents begin with higher k-level forecast rules than they did at time $t = 10$ – and in addition the average k-level continues to increase – with the result that the sacrifice ratio is considerably less than in the earlier disinflation.

There are two takeaways: First, in contrast to RE, the unified model of learning to forecast, appropriately calibrated, can provide reasonable realized sacrifice ratios. Second, in contrast to a model with a fixed profile of k-levels, the unified model provides a mechanism by which agents' sophistication can increase, thereby endogenously reducing the sacrifice ratio. Thus, after a sharp, and possibly painful deflation, the reputation earned by the CB makes the economy easier

[29] We set $\varphi = 0.8$, and the gain at $\phi = 0.1$, with the replicator coefficient $\alpha = 500$.

to stabilize.

6.2 Forward guidance

Within the context of a stylized new-Keynesian model, we consider the effects of announced policy in the form of a temporary interest-rate peg. This type of policy might most naturally be implemented when the economy is at the ZLB, but to remain close to the ad-hoc model, we simply assume that the economy is hit by a temporary negative demand shock of known duration, and policy makers combat the shock by announcing a temporary peg of the policy interest rate at a level below its target rate. Within this setting we can study the "forward guidance puzzle," first described by Giannoni, Patterson, Del Negro et al. (2015), and discussed in detail by Gibbs and McClung (2022). Informally, the puzzle states that New Keynesian RE models obtain implausibly large macroeconomic responses resulting from central bank forward guidance promises concerning future policy interest rates.

Our stylized model abstracts from the usual new-Keynesian Phillips curve by eliminating the dependence on expected inflation. This assumption has been adopted by a number of authors to engender tractability by reducing the dimension of the state variable: see Kocherlakota (2016), Williamson (2016) and Section 1 of Evans and McGough (2018). The model has the form

$$x_t = E_t x_{t+1} - \sigma^{-1}\left(i_t - E_t \pi_{t+1} - r\right) + v_t \tag{8}$$

$$\pi_t = b x_t, \tag{9}$$

where x_t is the output gap, π_t is inflation, r is the natural real rate, and $v_t = v^* + \varepsilon_t$, with $\varepsilon_t \sim \text{iid}(0, \sigma_\varepsilon^2)$, is a demand shock. Combining equations (8) and (9) gives

$$\pi_t = \theta \pi^e_{t+1} + (\theta - 1)r - (\theta - 1)i_t + b v_t, \tag{10}$$

where $\theta = 1 + b\sigma^{-1} > 1$.

For convenience, we now write the model in deviation from steady-state form. Let i^* and π^* be the inflation and nominal rate targets, respectively, and assume $i^* = \pi^* + r$ and $v^* = 0$. Then equation (10) becomes

$$\pi_t = \theta \pi^e_{t+1} - (\theta - 1)i_t + b v_t, \tag{11}$$

where, abusing notation to avoid clutter, π_t and i_t now represent period t deviations from their respective steady states. The policy rule, except when i_t is pegged at a fixed level, is given by $i_t = \varphi_\pi \pi^e_{t+1}$. Under this rule the economy's reduced form system is

$$\pi_t = \psi\pi^e_{t+1} + bv_t, \text{ where } \psi = \psi(\varphi_\pi) = \theta - \varphi_\pi(\theta - 1).$$

Since $\psi(1) = 1$ and $\psi' < 0$, it follows that the model is determinate provided $\varphi_\pi > 1$ is not too large. The REE, in this case, is $\pi_t = bv_t$.

We now consider the following experiment. Shut down the noise. At time $t = 0$ there is a temporary shock to steady-state demand $dv^* < 0$. Agents observe the shock, and are aware that it will last $M = M_1 + M_2$ periods. In period M_1 the CB announces a change in policy: it will peg the interest rate at $i^* + \Delta i$, where $\Delta i < 0$, for N periods, with $N \geq M_2$; then, in period $M_1 + N$ it will return to using the Taylor rule.

We first compute the REE. While $0 \leq t < M_1$, rational agents believe the Taylor rule will be implemented indefinitely, and that the demand shock will last M periods. If this were true, $\pi_n = 0$ for $n \geq M$. Backward inducting, we have

$$\pi_k = b\Delta v \left(\frac{1 - \psi^{M-k}}{1 - \psi}\right) \text{ for } 0 \leq k < M_1. \tag{12}$$

For $t \geq M_1$, the policy rule is changed and rational agents adjust accordingly. Reasoning as above, $\pi_n = 0$ for $n \geq M_1 + N$. Backward induction gives

$$\pi_k = \big(\chi\left(M_1 \leq k < M\right) b\Delta v - (\theta-1)\Delta i\big) \left(\frac{1 - \theta^{N-k}}{1 - \theta}\right) \text{ for } M_1 \leq k < N-1, \tag{13}$$

where $\chi(\cdot)$ in the Boolean indicator. Equations (12) and (13), together with $\pi_n = 0$ for $n \geq M_1 + N$, identify the REE under the temporary peg.

Turning now to the implementation of unified learning, the dynamics of the economy can be loosely written as

$$\pi_t = \gamma_t + \beta_t\pi^e_{t+1}, \tag{14}$$

though there is an important nuance: contemporaneous outcomes depend on (agents' perceptions of) the future path of model coefficients; and the entire path of coefficients is subject to change when policy announcements are made. However, at the time of the policy announcement, $t = M_1$, the model's time-path of parameters is fixed, and the dynamics are precisely captured by a system of the form (14). For simplicity of exposition, we discuss k-level forecasts of agents assuming $t \geq M_1$. A complete development is found in the online Appendix.

For the NK model (14), inflation is determined as follows:

$$\pi_t = \gamma_t + \beta_t \sum_{k\geq 0} \omega_t(k) E_t^k \pi_{t+1},$$

where $E_t^k \pi_{t+1}$ is the period t forecast of π_{t+1} made by a k-level agent. An expression for $E_t^k \pi_{t+1}$ can be derived using backward induction, starting with $E_t^1 \pi_{t+k} = \gamma_{t+k} + \beta_{t+k} a_t$, to obtain $E_t^k \pi_{t+1} = \sum_{n=1}^k \beta_t^{n-1} \gamma_{t+n} + \beta_t^k a_t$, where $\beta_t^k = \prod_{n=1}^k \beta_{t+n}$. The resulting actual law of motion is[30]

$$\pi_t = \gamma_t + \beta_t \sum_{k\geq 0} \sum_{n=1}^{k} \omega_t(k) \beta_t^{n-1} \gamma_{t+n} + \left(\beta_t \sum_{k\geq 0} \omega_t(k) \beta_t^k \right) a_t. \tag{15}$$

Equation (15) can be joined to the unified dynamic to simulate the economy with the adverse demand shock and forward-guidance policy specified above.

Figure 9 provides the results of a calibrated simulation. Here the slope of the Phillips curve is $b = 1.0$, and the response of the output gap to the real rate is $\sigma^{-1} = 1.0$. The gain is set at $\phi = 0.1$, and the replicator responsiveness is tuned with $\alpha = 50.0$.

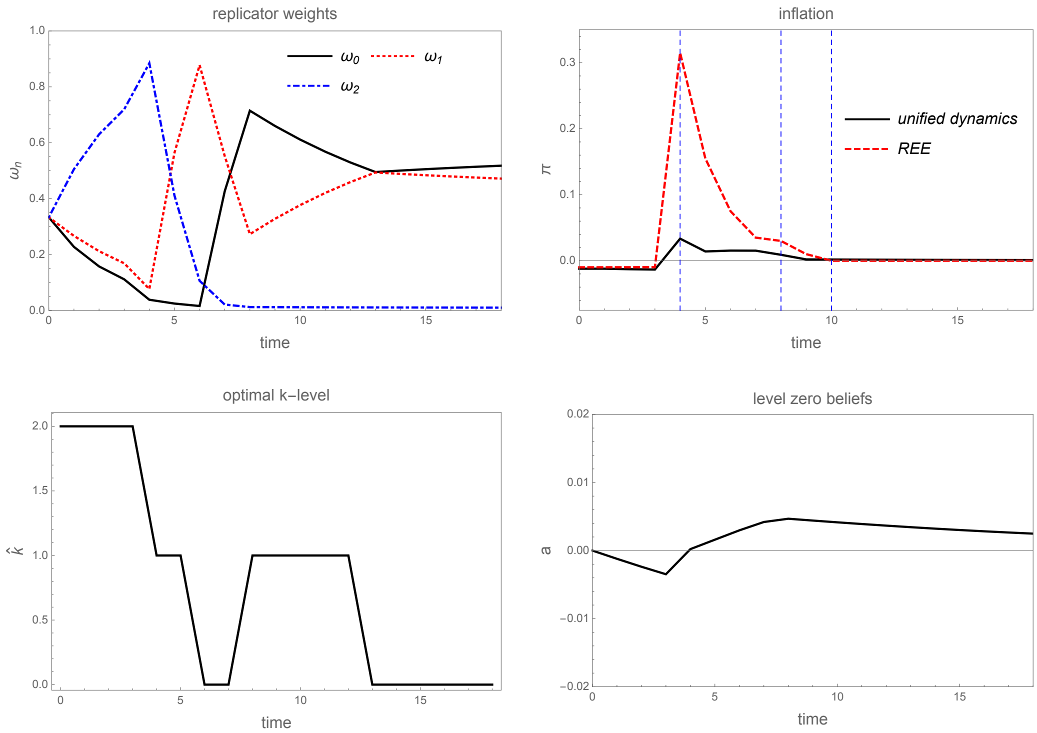

Figure 9: Forward guidance: $M_1 = 4, M_2 = 4, N = 6, \Delta v^* = -.005, \Delta i = -.01$

The specifics of the announced policy are as follows: in period $t = 0$, the economy is hit with a demand shock that results in a ceteris paribus 0.5% decline

[30] Details of this derivation are provided in an online Appendix A6.1.

in inflation each period for eight periods ($M_1+M_2 = 8$). In period $t = 4$ ($M_2 = 4$) the CB announces that it will hold the nominal rate at 1.0% below target for six periods ($N = 6$), after which it implements the active Taylor rule. This peg choice, when coupled with the negative demand shock, implies direct effects on inflation (for unchanged expectations) of a 0.5% increase in inflation in periods $t = 4, 5, 6, 7$ and a 1% increase in periods 8 and 9.

The NE panel illustrates the forward guidance puzzle under RE: an implausibly large impact response of inflation to the policy announcement, and it is straightforward to show that the impact effect is increasing in N. Under rationality, the proposed, and arguably moderate, policy results in an inflation spike of over 30%. Under unified learning, the response is much more modest, and in our view, more reasonable.

The modest response under unified learning derives explicitly from the replicator dynamics governing the proportion of k-level agents. At $t = 0$ we take initial beliefs to be at the prior REE value of $\pi = 0$, and we assume agents are uniformly distributed across levels $0, 1$ and 2. At the steady state REE there is no advantage to higher level forecasts; however, the presence of the demand shock changes the agents' perceptions of current and future values of γ_t. For the calibration under consideration, the optimal level is $k = 2$ for periods $t = 0, 1, 2, 3,$ and so $\omega_t(2)$ rises in periods $t = 1, 2, 3, 4$.

In period 4 the central bank announces its new policy. This again changes agents' perceptions of current and future values of γ_t, but this time perceptions of current and future values of β_t also change: for periods $t = 4, \ldots, 10$, agents perceive $\beta_t = \theta = 2$, and for $t > 10$ that $\beta_t = \psi = 0.5$. That β_t is larger than one for six periods implies that higher level forecasts iteratively magnify level zero beliefs, creating large forecast errors. For this reason, at the time of, and following the announcement, the optimal level falls to zero for a period before rebounding to level 1, resulting in the mitigated response of inflation to the announcement. The endogenous downward revisions to k provide an explanation for Farhi and Werning's (2019) result that low levels of deduction are required in order to resolve the foreword guidance puzzle across a range of macroeconomic models using level-k reasoning.

7 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 be a good description of 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.

Finally, our monetary policy applications to disinflation and the impact of interest rate forward guidance indicates the importance of extending the unified model to multivariate forward-looking settings, with potentially broad applications in macroeconomics and finance.

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.

Ball, Laurence, "What determines the sacrifice ratio?," in "Monetary policy," The University of Chicago Press, 1994, pp. 155–193.

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.

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.

Bullard, James, “Getting Ahead of U.S. Inflation: A Lesson from 1974 and 1983,” in “Federal Reserve Bank of St.Louis Regional Economist” Federal Reserve Bank of St. Louis 2022.

_ , “Reflections on the disinflationary methods of Poincaré and Thatcher,” in “Expectations in Dynamic Macroeconomic Models Annual Conference” Barcelona School of Economics Summer Forum 2022.

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.

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

Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Jane Ryngaert, “Do You Know that I Know that You Know. . . ? Higher-Order Beliefs in Survey Data,” *The Quarterly Journal of Economics*, 2021, *136* (3), 1387–1446.

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.

Crawford, Vincent P, Miguel A Costa-Gomes, and Nagore Iriberri, “Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications,” *Journal of Economic Literature*, 2013, *51* (1), 5–62.

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

Evans, George W., “The Fragility of Sunspots and Bubbles,” *Journal of Monetary Economics*, 1989, *23*, 297–317.

Evans, George W and Bruce McGough, “Interest-Rate Pegs in New Keynesian Models,” *Journal of Money, Credit and Banking*, 2018, *50* (5), 939–965.

_ **and Garey Ramey**, “Expectation calculation and macroeconomic dynamics,” *The American Economic Review*, 1992, pp. 207–224.

_ **and** _ , “Calculation, adaptation and rational expectations,” *Macroeconomic Dynamics*, 1998, *2* (2), 156–182.

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

_ **,** _ **, 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.

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.

Giannoni, Marc, Christina Patterson, Marco Del Negro et al., "The forward guidance puzzle," in "2015 Meeting Papers" number 1529 Society for Economic Dynamics 2015.

Gibbs, Christopher and Nigel McClung, "Does my model predict a forward guidance puzzle?," *Available at SSRN 3662025*, 2022.

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.

Greiner, Ben, "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE," *Journal of the Economic Science Association*, 2015, *1* (1), 114–125.

Guesnerie, Roger, "An exploration of the eductive justifications of the rational-expectations hypothesis," *The American Economic Review*, 1992, pp. 1254–1278.

_ , "Anchoring economic predictions in common knowledge," *Econometrica*, 2002, *70* (2), 439–480.

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.

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".

Kocherlakota, Narayana, "Neo-Fisherianism: A hopefully helpful analogy," *Kocherlakota's Thoughts on Policy blog post*, 2016.

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.

Muth, John F, "Rational expectations and the theory of price movements," *Econometrica: Journal of the Econometric Society*, 1961, pp. 315–335.

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

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

Sargent, Thomas J, "Rational expectations and inflation," in "Rational Expectations and Inflation," Princeton University Press, 2013.

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.

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

Williamson, Stephen, "More Neo-Fisher," *New Monetarist Economics blog post*, 2016.

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