

A Unified Model of Learning to Forecast*

George W. Evans

University of Oregon and University of St Andrews

Christopher G. Gibbs

Bruce McGough

University of Sydney

University of Oregon

November 1, 2021

ABSTRACT

We propose and experimentally test 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. The experimental results support these predictions. Our unified model is developed in a simple framework, but can clearly be extended to more general macroeconomic environments.

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

Key Words: expectations; adaptive learning; level-k reasoning; behavioral macroeconomics; 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, and the Expectations and Dynamic Macroeconomic Models Conference hosted by the University of Birmingham. 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, UNSW Business School, and The University of Sydney is gratefully acknowledged.

1 INTRODUCTION

The assumption of rational expectations (RE) continues to come under scrutiny in macroeconomics and finance models in which RE plays a central role. RE imposes strong assumptions on agents' knowledge and cognitive abilities that call into question the plausibility and robustness of some model predictions. This issue is particularly acute when studying the general equilibrium implications of structural change in RE models in which there are numerous empirical puzzles. Increasingly, modelers are turning to boundedly rational alternatives to RE such as adaptive learning (e.g. Evans, Honkapohja and Mitra, 2009 and Gibbs and Kulish, 2017), level- k reasoning (e.g. Angeletos and Lian, 2018 and Farhi and Werning, 2019), and behavioral models (e.g. Arifovic, Schmitt-Grohé and Uribe, 2018 and Goy, Hommes and Mavromatis, 2020) to attempt to resolve the puzzles. 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). 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), 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.¹

We assume agents agree on the level-0 forecast, are free to choose between k -level forecasts, and in general make heterogeneous choices. Over time the 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

¹ In the limit, the deductions comprise a key feature of the eductive learning framework of Guesnerie (1992) and Guesnerie (2002).

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.

To develop our unified approach, we consider a repeated static model, specifically the cobweb model as in Muth (1961), but allowing for possible positive feedback. This is the simplest model in which expectations affect the evolution of the economy; however, our approach can be extended to more general forward-looking environments like those associated with DSGE models. Our model is univariate and non-stochastic, and agents are differentiated by their reasoning level. The aggregate variable is determined by the expectations of the agents: $y_t = \gamma + \beta \hat{E}_{t-1} y_t$, where $\hat{E}_{t-1} y_t$ is the average of individual forecasts. The setup can also be thought of as a repeated beauty contest or guess-the-average game, where β is the feedback from expectations to economic outcomes. Thus we consider both positive and negative expectational feedback cases, $\beta > 0$ and $\beta < 0$.

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 the linear cobweb model under adaptive learning, educative learning, and RE are of course 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 by agents making deductions via educative reasoning only obtains if $|\beta| < 1$.²

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 the through the proportions of agents using different k -levels. Quick convergence can result in the coexistence of high-level and low-level reasoners for extended periods. Once the market has approximately converged, all level- k forecasts can provide almost the same pre-

²The educative approach is based on strong common knowledge assumptions. See Guesnerie (1992) for application of this approach to the cobweb model.

diction, 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 may 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 highest level reasoners may make large forecast errors. In this case, some mass 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.

We adapt the experimental design of Bao and Duffy (2016) to test the key elements of the unified model in an LtFE. We place laboratory subjects into a computer based market that nests the cobweb model. The market price depends on one-period ahead 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 peri-

ods, 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, behavior that is sometimes argued to be implausible when level-k reasoning is adapted to more complex macroeconomic environments as in García-Schmidt and Woodford (2019) and Angeletos and Sastry (2021).

We also find evidence for 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 prevalence 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.

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 Bray and Savin (1986), Marcet and Sargent (1989), and Evans and Honkapohja (2001) to ground the level-0 forecast assumptions and analysis. For behavioral heterogeneity assumptions, we draw heavily on the work of Brock and Hommes (1997), De Grauwe (2012) and Hommes (2013), which all consider ex ante homogeneous agents that select from a menu of forecast rules leading to ex post heterogeneity in a variety of macroeconomic settings. We differ from these treat-

ments by considering menus with arbitrarily many forecast types. Agents in our framework are free to scale up or down their depth of reasoning as they see fit. Unlike this 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 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 . 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 (2019), 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. (2019), 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} .³ For $k \in \mathbb{N}$, the proportion ω_k of agents of type k (i.e. having sophistication k) is referred to as the *weight* associated with agent-type k . The distribution of agents across types is summarized by a *weight system* $\omega = \{\omega_0, \dots, \omega_M\}$, 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.⁴ 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 $E_{t-1}^k y_t$, where higher k indicates greater sophistication. The aggregate y_t is determined as

$$y_t = \gamma + \beta \sum_{k=0}^M \omega_k E_{t-1}^k y_t \equiv \gamma + \beta \sum_k \omega_k E_{t-1}^k y_t, \quad (1)$$

³To be clear, \mathbb{N} is assumed to include zero.

⁴It will sometimes be convenient to interpret a weight system as a sequence with only finitely many nonzero terms.

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).⁵ Agents with *level-0 beliefs* hold a common prior and form their forecasts accordingly as $E_{t-1}^0 y_t = a$. 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 $k - 1$:

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

Note that for $k \geq 1$ agents are assumed to know β and γ .⁶

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

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

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

⁵It 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, Ho and Chong (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.

⁶This 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 $k \geq 1$.

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.⁷ Specifically, a weight system ω is taken as fixed and level-0 forecasts $E_{t-1}^0 y_t \equiv a_{t-1}$ are assumed to evolve over time in response to observed outcomes. The system under adaptive dynamics is given by

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

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} , for $i \geq 1$.⁸ 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

⁷We 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.

⁸If 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.

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

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

where the left-most “min” is used to break ties.⁹

Next, assume the presence of a rate function $r : [0, \infty) \rightarrow [\delta, 1)$ with $\delta \geq 0$ satisfying $r' > 0$.¹⁰ 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

$$y_t = \gamma + \beta \sum_{k \geq 0} \omega_{kt} E_{t-1}^k y_t$$

$$\omega_{it+1} = \begin{cases} \omega_{it} + \sum_{j \neq \hat{k}(y_t)} r(|E_{t-1}^j y_t - y_t|) \omega_{jt} & \text{if } i = \hat{k}(y_t) \\ (1 - r(|E_{t-1}^i y_t - y_t|)) \omega_{it} & \text{else} \end{cases} \quad (5)$$

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

⁹Note that the existence of the argmin is guaranteed by the fact that if $|\beta| < 1$ then $E_{t-1}^k y_t \rightarrow 0$ as $k \rightarrow \infty$, and if $|\beta| > 1$ then $E_{t-1}^k y_t \rightarrow \infty$ as $k \rightarrow \infty$.

¹⁰An 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.

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

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

where $\delta_r \in \{0, 1\}$ indicates whether the replicator dynamic is operable. We note that while the adaptive learning 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 $|y_t| \rightarrow \infty$ 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.

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 learning 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.¹¹

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

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

¹¹Proofs of all theorems and propositions are found in the online appendix A1.

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

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

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

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

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^n s are pairwise disjoint. The set Ω inherits a natural topology, sometimes called the *final topology*, from the relative topologies on the Δ^n s: $W \subset \Omega$ is open if and only if $W = \dot{\cup}_n W_n$, with $W_n \subset \Delta^n$ open in Δ^n .¹²

Using this notation, and given $\beta \in \mathbb{R}$, we may define $\psi_\beta : \Omega \rightarrow \mathbb{R}$ 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 learning.

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

1. *If $|\beta| < 1$ then the model is stable: $y_t \rightarrow \bar{y}$.*
2. *If $\beta > 1$ then the model is unstable: $|y_t| \rightarrow \infty$.*
3. *If $\beta < -1$ then ψ_β is surjective, and*

¹²The final topology on a disjoint union of topological spaces is the *direct limit* topology induced by the inclusion maps $\Delta^n \hookrightarrow \Omega$.

- (a) If $\psi_\beta(\omega) > 1$ then the model is unstable: $|y_t| \rightarrow \infty$.
- (b) If $1 - 2\phi^{-1} < \psi_\beta(\omega) < 1$ then the model is stable: $y_t \rightarrow \bar{y}$.
- (c) If $\psi_\beta(\omega) < 1 - 2\phi^{-1}$ then model is unstable: $|y_t| \rightarrow \infty$.
- (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: $y_t \rightarrow \bar{y}$. ii) If $\omega \in \Omega_u$ then the model is unstable: $|y_t| \rightarrow \infty$. 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 learning dynamics may be written

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

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

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

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

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

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.

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

$$\hat{k}(a) = \min \arg \min_{k \in \mathbb{N}} |\beta^k a - \beta \sum_k \omega_k a|, \quad (7)$$

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

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

1. *Suppose $0 < \beta < 1$.*

$$(a) \ K \rightarrow \infty \implies \hat{k} \rightarrow \infty \text{ and } \hat{k}/K \rightarrow 0.$$

$$(b) \ \beta \rightarrow 1^- \implies \hat{k} \rightarrow \begin{cases} \frac{K}{2} + 1 & \text{if } K \text{ is even} \\ \frac{K+1}{2} & \text{if } K \text{ is odd} \end{cases}$$

$$(c) \ \beta \rightarrow 0^+ \implies \hat{k} \rightarrow \begin{cases} 1 & \text{if } K = 1 \\ 2 & \text{if } K \geq 2 \end{cases}$$

2. *Suppose $-1 < \beta < 0$.*

$$(a) \ K \rightarrow \infty \implies \hat{k} \rightarrow \infty \text{ and } \hat{k}/K \rightarrow 0.$$

$$(b) \ \beta \rightarrow 0^- \implies \hat{k} \rightarrow \begin{cases} 1 & \text{if } K = 1 \\ 3 & \text{if } K \geq 2 \end{cases}$$

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

3. *Suppose $\beta < -1$*

$$(a) \ K \rightarrow \infty \implies \hat{k} \rightarrow \infty \text{ and } \hat{k}/K \rightarrow 1$$

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

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

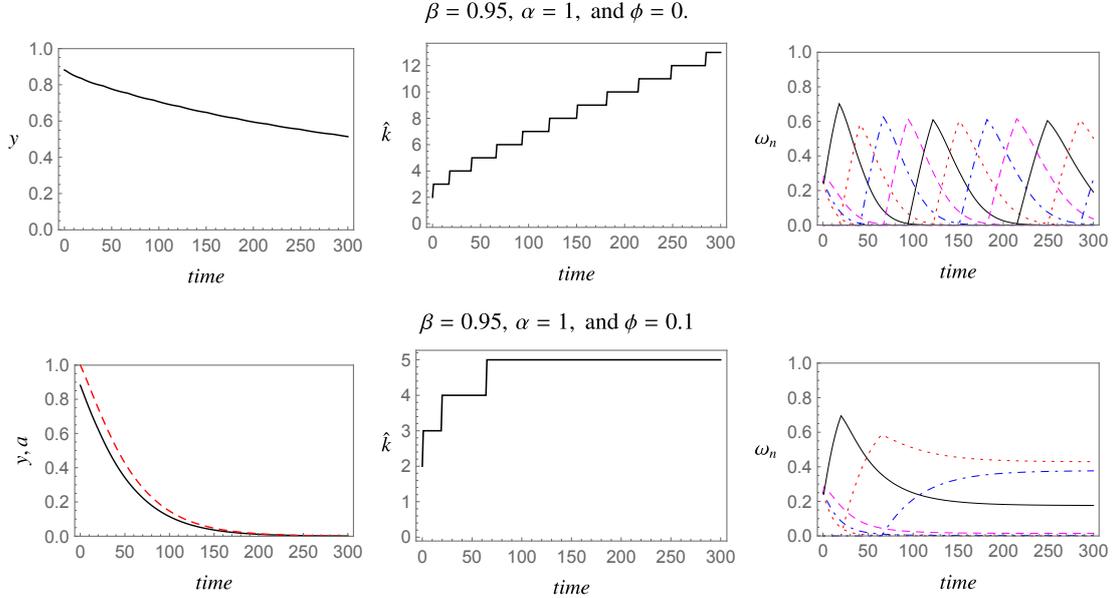
Although Proposition 1 examines only the specific case of uniform weights, it reveals the contrasting results for the optimal choice of k in negative and positive feedback cases. Item 1, which concerns positive feedback, shows that for large K , i.e. if agents are distributed across a wide range of sophistication, an approximately optimal forecast can be obtained with level- k beliefs with relatively low k . Items 2 and 3 reveal the challenges of choosing the optimal k when feedback is negative. For example, if β is slightly larger than -1 then \hat{k} is very large, which, in practice, would be cognitively taxing to determine. On the other hand, if β is slightly smaller than -1 then \hat{k} is very small: indeed, $\hat{k} \in \{0, 1\}$, with the specific value determined by the *aggregate parity*, which can be viewed as an aggregate measure of optimism and pessimism.

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

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 panels correspond to replicator dynamics ($\phi = 0$) and bottom panels to unified dynamics: we omit results associated with adaptive learning as they simply show monotonic convergence of a and y to \bar{y} .

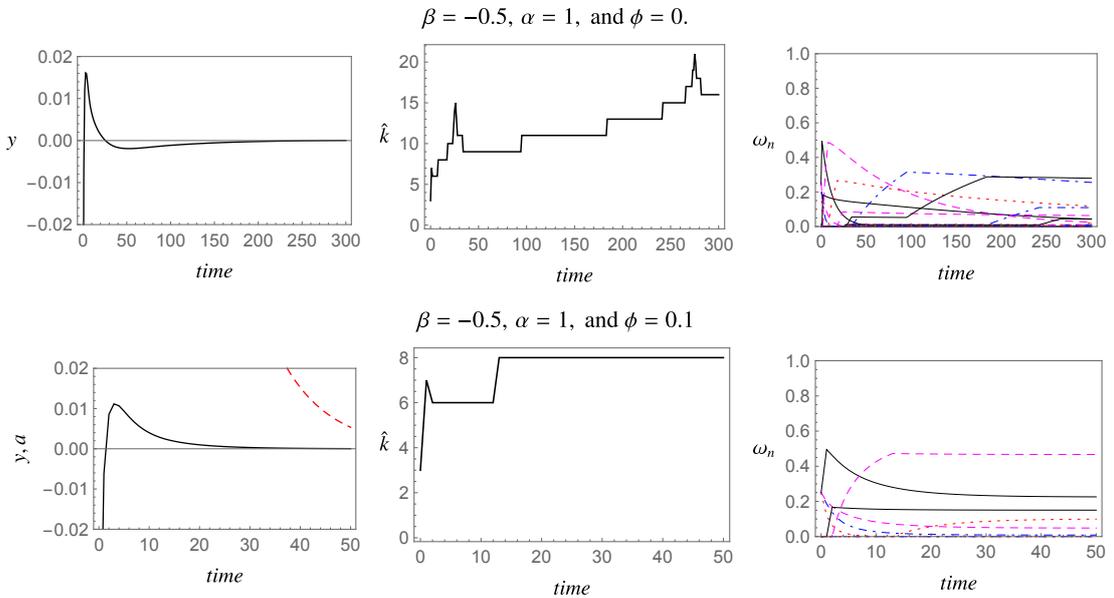
Under replicator dynamics, y exhibits monotone convergence to \bar{y} , as the weight distribution shifts to higher k -level forecasts. However, convergence of y to zero is slow relative to adaptive learning (not shown): Under the replicator, at the end of 300 periods, y has fallen less than half of the distance from the initial value to \bar{y} , whereas under adaptive learning with $\phi = 0.1$ y falls by over 90% in the same number of periods. However, the unified dynamics converges fastest: as seen in the lower left panel, after 300 periods y is indistinguishable from REE.

The right panel of Figure 1 provides the dynamics of agents' weights. As indicated above, $\omega_{n0} = 1/4$ 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. As the replicator adds higher forecast levels, the associated paths are graphically identified in an analogous fashion by repeating the styles mod four. We follow this convention throughout this section. As predicted by our theoretical results, we observe lower-level forecasts to gradually

fall out of favor and be replaced by higher-level forecasts.

Figure 1 also gives the results when both adaptive learning and replicator dynamics are operational. Here $\beta = 0.95$, $\phi = 0.1$ and $\alpha = 1$. The bottom left panel includes the time paths of both y (solid) and a (dashed). Convergence is now much faster, and also faster than the purely adaptive learning case. This reflects the combined role of adaptive learning in shifting the level-0 forecast over time toward \bar{y} and shifting weight toward higher values of k . Note that the weight associated to level k forecasts is increasing if and only if $k = \hat{k}$. For example, we see that around $t = 20$, weight begins to shift to level-4 forecasts, and to level-5 forecasts around $t = 65$. The optimal k appears to stall out at $\hat{k} = 5$ because, as the estimate $a_t \rightarrow 0$, higher-level forecasts provide limited to no additional value. Indeed, we have found in other simulations that if the gain is reduced, so that $a_t \rightarrow 0$ more slowly, \hat{k} rises further before appearing to plateau.

Figure 2: Simulated dynamics with negative feedback.

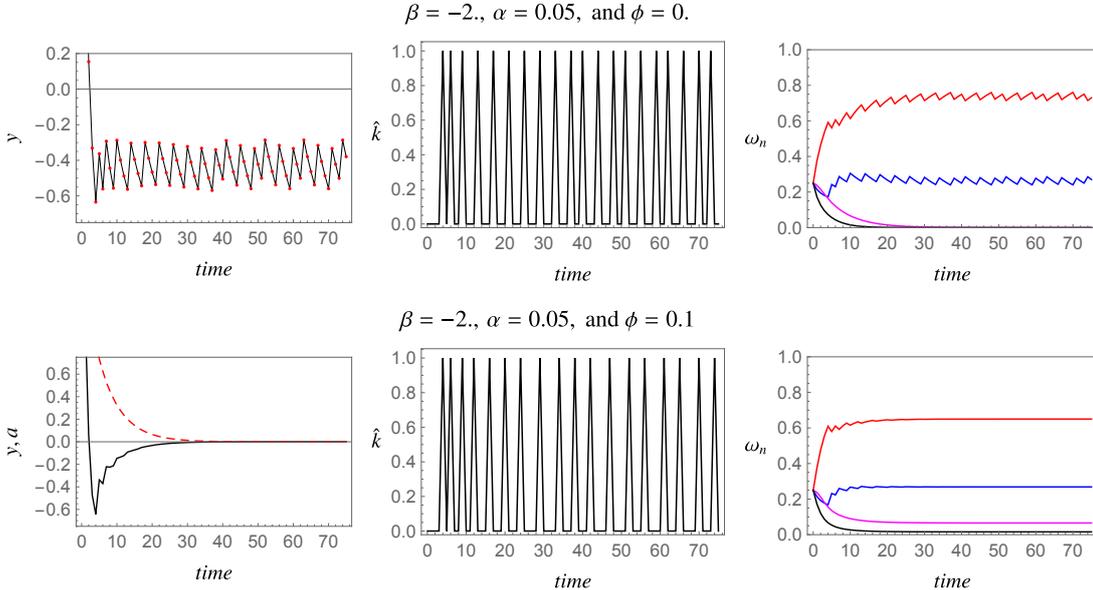


Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: red dotted, blue dash-dot, dashed magenta and solid black, respectively.

We now turn to the negative feedback case, with $-1 < \beta < 0$. Again, the results associated with adaptive learning are unexceptional, and indeed, for small ϕ , result in monotonic convergence of a and y . Turning to the replicator dynamics, Figure 2 provides the results for $\beta = -0.5$. First note the non-monotonic behavior of y : the left panel, scaled differently in this figure, shows oscillatory convergence of y induced by the negative feedback. The behavior of \hat{k} reflects these oscillations: when y crosses zero, \hat{k} rises sharply to drive down (in magnitude) the optimal forecast $\beta^{\hat{k}}$. There are, in fact, three times that y crosses zero during the observation period, at $t = 2, 26$, and 276 , which are associated

with spikes in \hat{k} , particularly visible in the latter two cases. Finally, we note that by Theorem 2, $\hat{k} \rightarrow \infty$. However, unlike the positive feedback case, here this convergence is not monotone. Figure 2 also gives the results when both adaptive and replicator dynamics are operational. Because adaptive learning drives level-0 forecasts to zero there is faster convergence, with weaker oscillatory behavior, than in the replicator-only case.

Figure 3: Simulated dynamics with large negative feedback.



Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). In the right panel, which plots the time paths of the weights, the higher oscillating path corresponds to ω_0 and the lower one to ω_1 .

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 $n \geq 2$. For viewing convenience, the right panel, which plots the time paths of these weights, does not employ the style convention used above; instead, the higher oscillating path corresponds to ω_0 and the lower one to ω_1 . The evolution of y appears to converge to an 11-cycle, which, we observe, is not centered at zero.¹³ The bottom panel of Figure 3 exhibits the corresponding simulation with unified dynamics in which $\phi = 0.1$. The addition of adaptive learning dynamics

¹³We find numerically that there are at least two stable 11-cycles.

pushes level-0 expectations towards zero, which when combined with replicator dynamics leads to rapid convergence to the REE.¹⁴

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 $E_{t-1}^k y_t$.

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.¹⁵ 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.

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 .

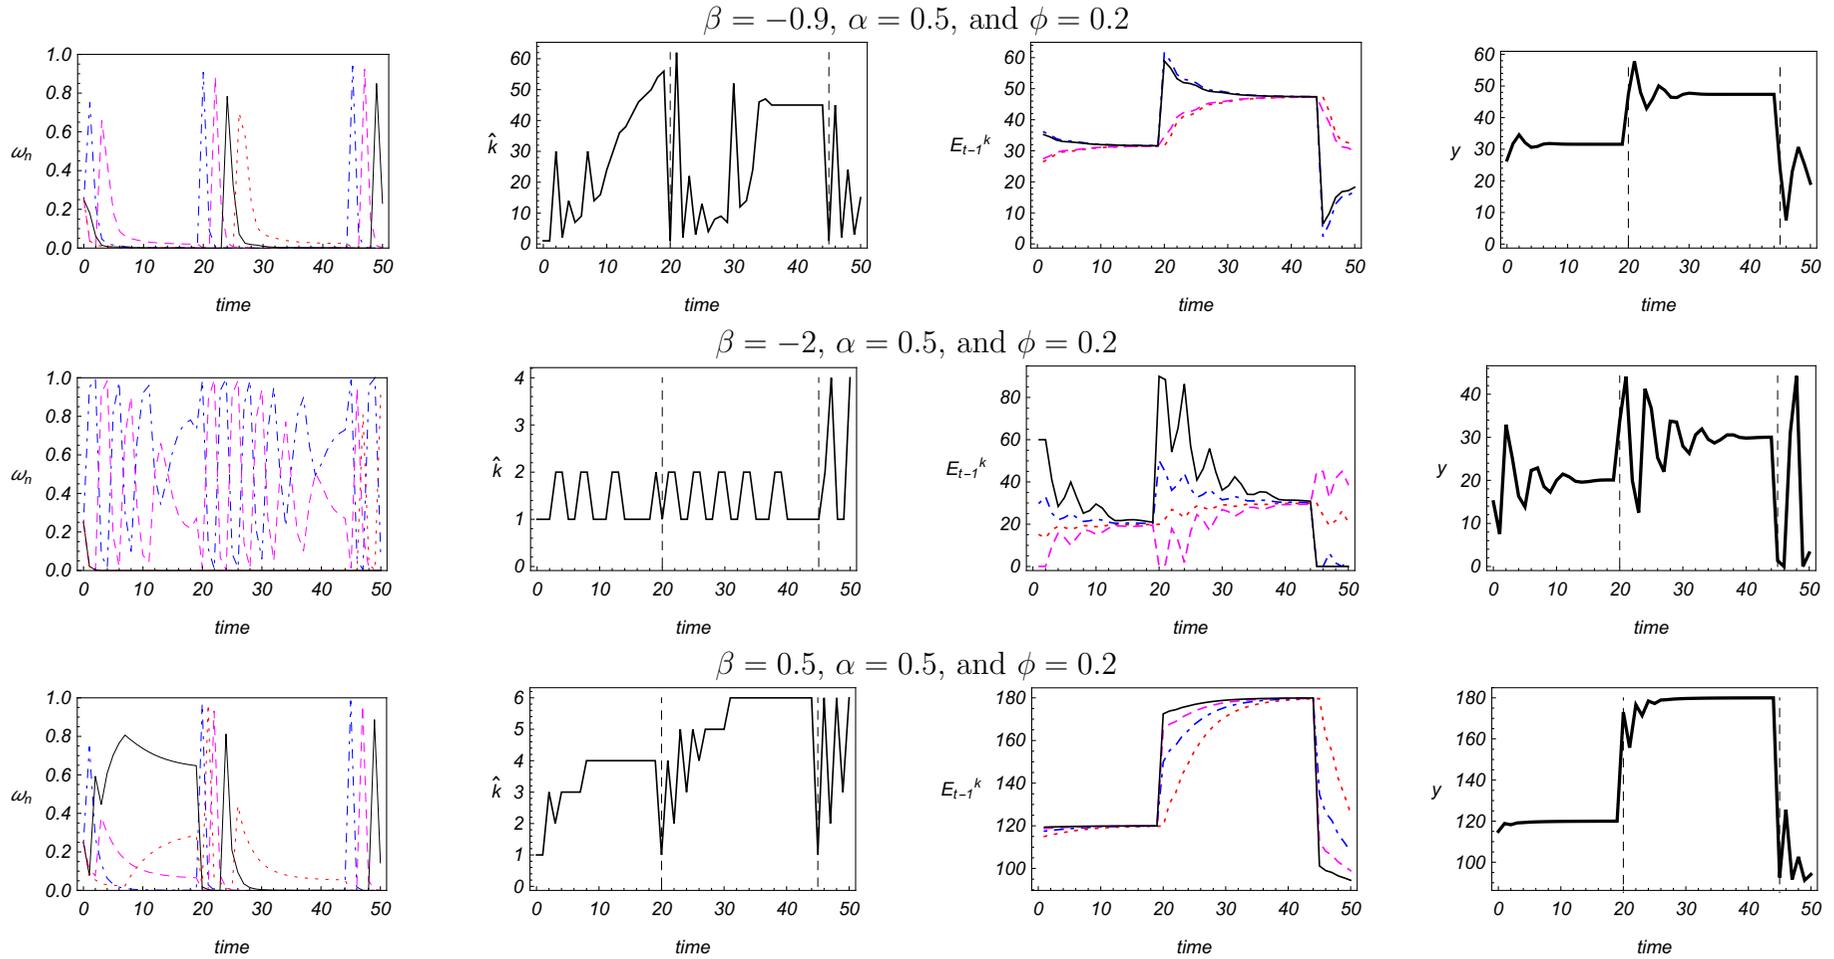
Starting with the $\beta = 0.5$ simulation, we note three additional features of the unified dynamics. First, despite the fact that $y_t = \bar{y}$ for many periods prior to the announcements, the model does not predict instantaneous convergence to the new REE in these periods. In other words, convergence 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. The similarity among all the predictors lowers the incentives to revise

¹⁴Larger values of α may lead to explosive replicator dynamics; and, we have found, for our calibrations, that the corresponding unified dynamics continue to induce convergence.

¹⁵The 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. Generalizing the unified dynamic to models with timing $E_t y_{t+1}$ is of considerable interest and under current investigation. In this environment, how soon the information is incorporated into forecasts depends on reasoning level.

one's depth of reasoning higher 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 as predicted by Proposition 1 (see the second plot in the last 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 plots in the last 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 in every period (see the first plot and recall that the proportion using $k > 3$ is not shown).

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, $E_{t-1}^k y_t$, use the same style format.

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.

These same dynamics are repeated for wide range of parameters when $|\beta| < 1$. The consistency of the result is owed to the relatively quick convergence of y_t to \bar{y} when $|\beta| < 1$. Because there is less incentive to continue revising one's depth of reasoning as y_t approaches \bar{y} , there is always a mass of low-level reasoners when the announcements occur, which triggers the dynamics shown in Figure 4. However, because not all high-level reasoners will revise their depth of reasoning in response to the forecast error in an announcement period, 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. The parameters we select show a case where 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)

Table 1: Experimental Treatments

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

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.¹⁶

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, by everyone observing market price dynamics, there is an obvious level-0 forecast from which to make level-k deductions. Figure 4 shows that individual forecasts are heterogeneous and should diverge dramatically from what we observe in the periods prior to the announcement, which allows us to precisely characterize whether individual forecasts coincide with the unified model.

Hypothesis 3 (Replicator Dynamics): *In response to losses, some participants revise their level of reasoning to the current optimal predictor.*

¹⁶Bao and Duffy (2016) note that the T2 treatment does not satisfy educative 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 educative 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 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.*

Finally, our simulations suggest that for treatments T1 and T3, a change in structure can lead to a temporary reduction in the average depth of reasoning.

4.1 EXPERIMENT DESCRIPTION

The experiment used a computer based market programmed in oTree.¹⁷ 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.

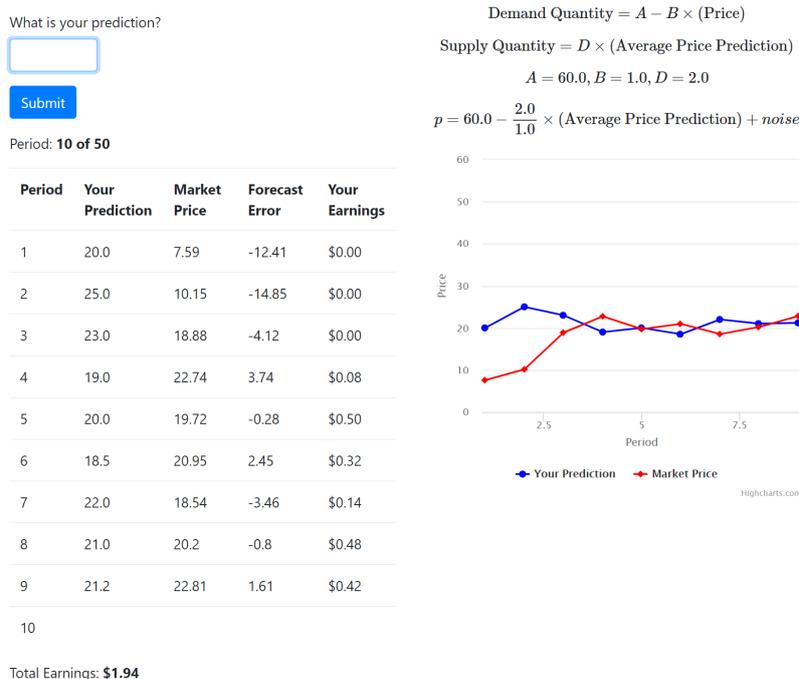
Participants were given slightly different stories about the market environment in the positive (T3) and the negative feedback cases (T1 and T2). In the latter, participants were told that the market follows the normal cobweb setup of perfect competition among firms that face convex costs of production 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:

¹⁷See Chen, Schonger and Wickens (2016) for documentation.

Figure 5: Screenshot of experimental market GUI

Market Game



Consider the case where $A = 60, B = 2, D = 1$ and $\text{noise} = 0$. If we substitute these numbers into this equation

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

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.¹⁸ 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 [A6](#).

Figure 5 shows the graphic user interface (GUI) that participants interacted with during the experiment. The market information is shown in the top right

¹⁸Four 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.

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.

The payoff function for the participant’s predictions is

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

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.¹⁹ 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.

Each participant played 50 rounds. 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

¹⁹Ethics requirements placed on the study mandated that participant payments were on average \$15 AUD per hour.

Table 2: Summary statistics

Treatments	Markets (62)	Participants (372)	Treatment Values		Payments		Time Use (min)	
			Feedback	Announcements	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.

each session. Participant no-shows account for the different number of markets across the treatments.²⁰ 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.²¹ 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.

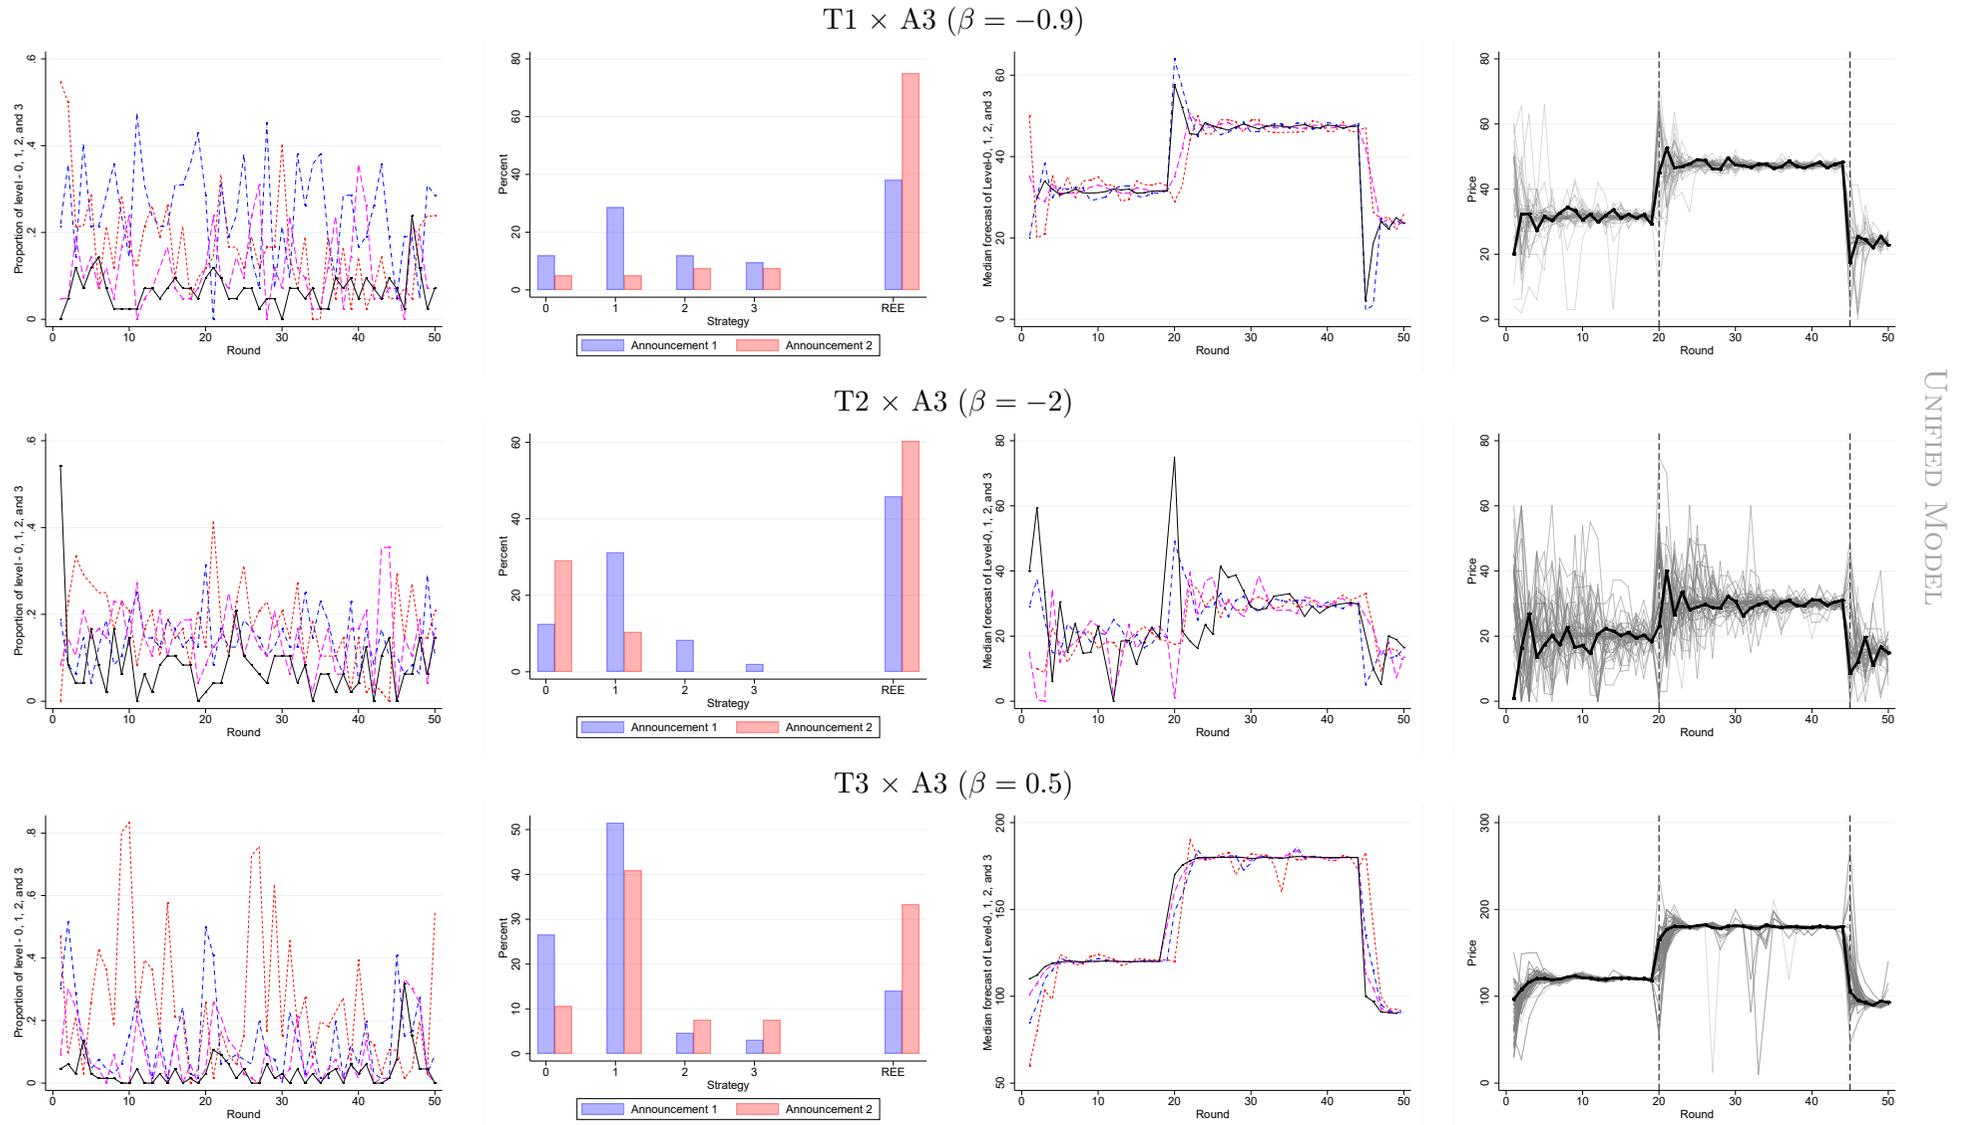
5.1 CONVERGENCE RESULTS

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.

²⁰Subjects were recruited using ORSEE (see Greiner, 2015).

²¹The 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



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.

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

The second metric we use to assess convergence is the mean difference in the

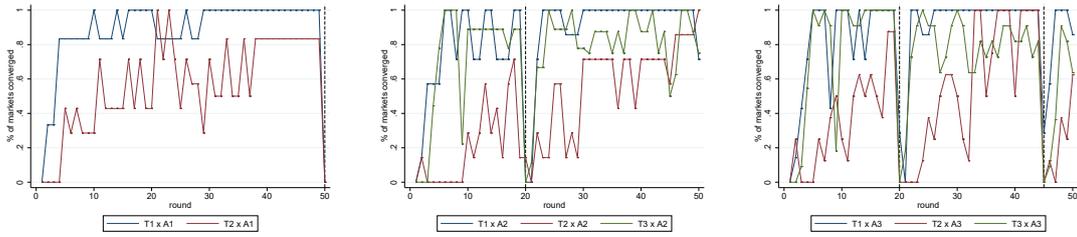
market price from steady state over the same intervals used for the first metric. Columns 5 - 7 of Table 3 show the mean difference and the t-statistics for a test of the null hypothesis that the mean difference is less-than-or-equal to 3. Bolded values indicate a one-sided rejection of the null hypothesis 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 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.

Figure 7 provides further visualization of the difference in convergence across feedback treatments. The plots show the proportion of markets that recorded two consecutive periods within the plus or minus three band for a rolling window. The T1×A1 and T2×A1 treatments are particularly informative on convergence here. These treatments were designed to closely replicate the main treatment of Bao and Duffy (2016). We replicate their main results by finding relatively quick and maintained market convergence for the T1 treatment and slow or no convergence for the T2 treatments. The T2 treatment also provided the greatest heterogeneity for the speed of convergence with some markets never converging – a phenomena not observed for any market in the T1 and T3 treatments – and some converging quite quickly.²² The heterogeneous outcomes for the T2 treatments

²²One case of non-convergence is particularly noteworthy. One of the T2 x A1 experiments collapsed completely at period 36. One participant was frustrated with the lack of earnings and requested to leave the experiment. All participants in this market were paid their show-up fee and earnings up to period 36 and then dismissed from the experiment.

Figure 7: Convergence of price to REE in experimental markets



Notes: The plots show the percentage of markets that have converged by the round indicated for different treatments. Convergence is defined as being within ± 3 of the steady state price for two consecutive quarters on a rolling basis.

are predicted by the unified model. Small changes in initial conditions to any number of different parameters of the model can lead to coordination on the REE or to a market that completely destabilizes. Further, the unified model correctly predicts that all T1 and T3 treatments should robustly converge.

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.²³ 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

²³The results are robust to reasonable changes in the definition of level-0 forecast. In the online appendix A3, 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 A2.1, which illustrates further how the classification works in practice.

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 hypothesis, we estimate the following regression model:

$$d_{i,r} = \alpha_i + \omega_r + \sum_{k=0}^3 \beta_k I(k=1)_{i,r} + \sum_{j=0}^3 \gamma_j I(j=1)_{i,r} \times I(r=Ann)_r + \epsilon_{i,r} \quad (8)$$

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

Table 4: Classifying participant’s forecasts as Level-k

Within ± 3 of Level-k in announcement rounds				Differences in deliberation time (seconds)		
	1	20/50	45	Variable	(1)	(2)
Total Classified	47.3% [33.8% , 56.9%]	64.4% [52.6% , 71.6%]	66.0% [48.1% , 70.5%]	Level 0	-5.84 (0.854)	-1.31 (0.565)
Level-0	14.8% [11.0% , 15.1%]	6.6% [4.31% , 8.05%]	5.1% [4.49% , 7.05%]	Level-1	-4.89 (0.950)	-0.90 (0.701)
Level-1	7.3% [6.45% , 8.60%]	24.1% [20.7% , 26.7%]	14.1% [12.2% , 14.1%]	Level-2	-3.96 (1.210)	-1.13 (0.851)
Level-2	6.5% [1.88% , 6.45%]	5.5% [4.60% , 5.75%]	3.8% [1.92% , 3.85%]	Level-3	-3.82 (1.370)	0.21 (1.113)
Level-3	3.2% [1.11% , 11.3%]	3.4% [2.87% , 4.02%]	4.5% [3.85% , 5.13%]	Level-0 x Ann	45.25 (8.763)	2.44 (6.021)
REE	15.6% [13.4% , 15.6%]	24.7% [20.1% , 27.0%]	38% [25.6% , 40.4%]	Level-1 x Ann	43.12 (4.712)	12.25 (4.554)
				Level-2 x Ann	59.58 (8.806)	12.85 (8.390)
				Level-3 x Ann	62.83 (11.80)	22.31 (8.392)
N	372	348	156	Cons	39.5 (0.457)	112.68 (4.206)
Hypothesis tests of deliberation time regressions				Individual FE	yes	yes
				Round FE	no	yes
$H_0 : \text{Level-0} - \text{Level-3} = 0$				R-squared	0.027	0.253
$H_0 : (\text{Level-0} \times \text{Ann}) - (\text{Level-3} \times \text{Ann}) = 0$				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.

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 8 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.²⁴ 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

²⁴In the online appendix A4, we provide more evidence that subjects make oscillating deductions as the market converges, consistent with level-k reasoning.

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 8 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. In the interest of space, the survey results are discussed in the online appendix A5.

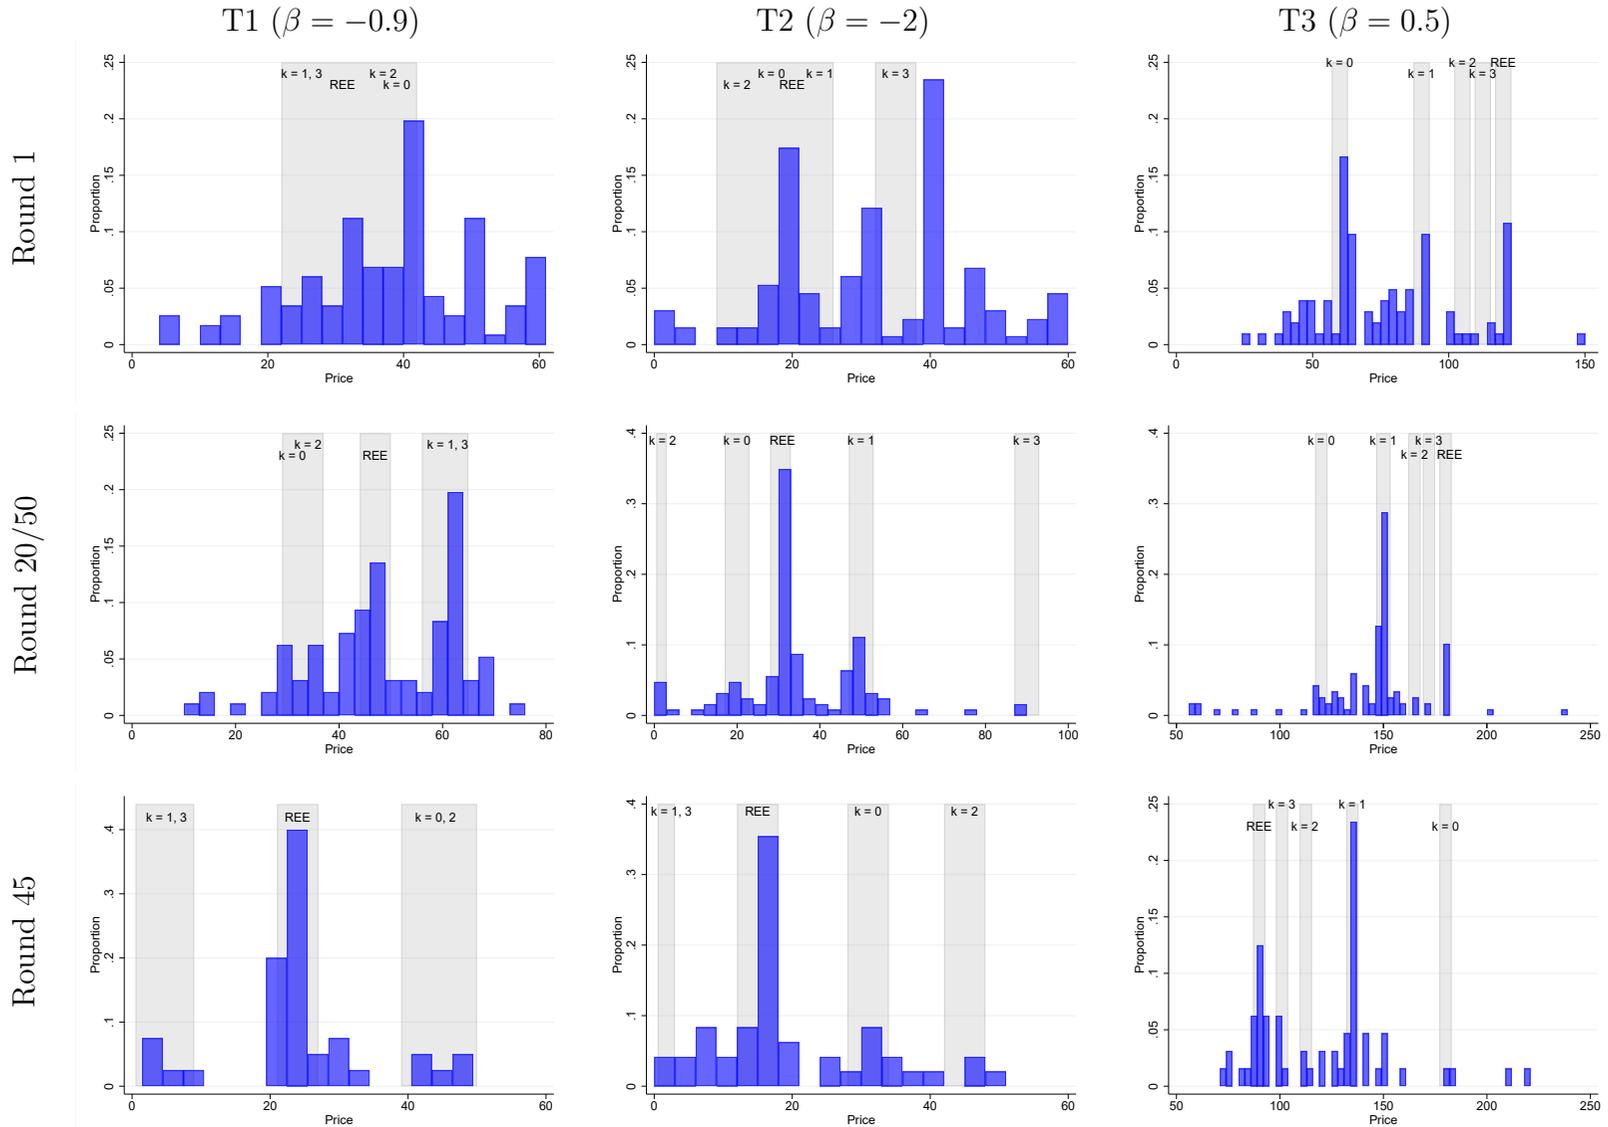
Finally, we return to Figure 6 that summarizes the overall dynamics we observe in the data. Here we do not make use of cutoffs for the classifications and instead use all of the data to summarize overall observed behavior. To do so, each individual forecast is classified as level-0, 1, 2, or 3 based on the whichever forecast it is closest to measured by absolute error.²⁵ 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 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.

²⁵Although 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 cutoffs, 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 A2.

Figure 8: 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.

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 from 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.

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 x 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 con-

Table 5: Revisions and loss

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

stant. 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.

5.4 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 level- k predicted price forecasts in the histograms shown in Figure 8. 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 \times 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.

We speculate that the high proportion of REE forecasts observed in the T2 \times 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, it's 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 CONCLUSION

The union of behavioral heterogeneity, adaptive learning, and level- k reasoning brings together three behavioral assumptions that enjoy wide experimental sup-

port. 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. Our 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. This in turns 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**, “Planar Beauty Contests,” *UNSW Business School Research Paper Forthcoming*, 2019.
- Arifovic, Jasmina, Stephanie Schmitt-Grohé, and Martin Uribe**, “Learning to live in a liquidity trap,” *Journal of Economic Dynamics and Control*, 2018, 89, 120–136.
- Bao, Te and John Duffy**, “Adaptive versus educative 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.
- 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.
- 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 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.
- 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”.
- 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.
- 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.