

A Unified Model of Learning to Forecast*

George W. Evans

University of Oregon and University of St Andrews

Christopher G. Gibbs

University of Sydney

Bruce McGough

University of Oregon

January 10, 2024

ABSTRACT

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

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

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

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

1 INTRODUCTION

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

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

The equilibrium nature of RE is seen most clearly in the simple guess-the-average game. Subjects pick a number between 0 and 100, with the winning number being closest to $2/3$ of the average guess. A subject who treats 50 as a focal point – the mean of a random guess from $[0,100]$ – might then choose the “level-1” guess of $(2/3) \times 50$. However a subject who thinks that other subjects make the level-1 guess, may make the level-2 guess of $(2/3)^2 \times 50$, etc. The unique Nash equilibrium guess is zero, but in one-shot games, even for grandmaster chess players, average guesses are typically over 30, with winning guesses over 20: see Nagel, Bühren, and Frank (2017). In many RE models, including Muth (1961), Lucas (1972), and the New Keynesian model, this type of strategic uncertainty is implicitly present but often ignored.

We seek to unify key elements of alternative bounded-rationality approaches to strategic uncertainty by marrying AL and level-k reasoning

¹In our formal model, we assume agents have full information about the structure of the economy. It is possible to retain RE while relaxing the full information assumption, e.g. Bianchi, Lettau, and Ludvigson (2022) and Bianchi and Melosi (2014) consider models in which agents are rational but face uncertainty about observed regime change. In our macroeconomic policy example in Section 6, we extend our approach to a setting with uncertainty about the duration of an observed policy regime.

within a single heterogeneous-expectations behavioral model. Adaptive learning and heterogeneous expectations capture the 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).² Level- k reasoning provides a way to model general equilibrium implications of forward-looking boundedly rational expectations, and has wide experimental support: see Nagel (1995), Duffy and Nagel (1997), Ho, Camerer, and Weigelt (1998), Bosch-Domenech, Montalvo, Nagel, and Satorra (2002), Costa-Gomes and Crawford (2006), Nagel, Bühren, and Frank (2017) and Mauersberger and Nagel (2018). This literature also shows that in repeated “guess-the-average” games, a special case of our univariate model, agents tend to shift over time to higher level- k forecasts.

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

Central to our approach is that agents (i) have knowledge about the economic structure – specifically about how outcomes depend on the expectations of other agents – but (ii) cannot directly observe the expectations of other agents. Rational agents may understand that there is an RE equilibrium (REE), yet refrain from holding expectations consistent with the REE because they doubt that other agents will hold RE.³ The validity of these doubts is well-established in the above-cited experimental literature.

Of course, rational agents might well consider the possibility that other agents have heterogeneous expectations, in particular that other agents

²LtFE are laboratory experiments in which the sole or principal task of the subject is to make forecasts of key economic variables.

³The eductive learning literature, e.g. Guesnerie (1992), Guesnerie (2002), emphasizes that both structural restrictions and strong higher-order common knowledge assumptions would be needed for fully rational agents to coordinate on the REE. Eductive reasoning more broadly refers to using knowledge of the economic structure to make inferences about the possible expectations held by other rational agents.

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

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

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

To summarize: our bounded rationality model, which we call *unified dynamics*, includes three elements:

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

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

We study the unified model first in the univariate set-up of Muth (1961). After deriving sharp analytical results and examining simulations that illustrate the model's implications for different types of expectational feedback,

we take the model to the laboratory and test its core predictions using a standard experimental design. We then extend our framework to the New Keynesian model. We show that the unified model justifies low level-k assumptions adopted in prominent papers such as Angeletos and Lian (2018) and Farhi and Werning (2019).

There are other approaches in the literature. One strand assumes fixed proportions of agents that differ in their sophistication. The simplest cases include two categories: unsophisticated agents and fully rational agents who take into account the proportion of nonrational agents. Gali and Gertler (1999) and Jackson (2005) consider inflation dynamics when a fixed proportion of agents follow naive rule-of-thumb forecasts, while the other agents are fully rational. Mokhtarzadeh and Petersen (2021) explore a monetary model in which a proportion of agents have expectations aligned with central bank forecasts, while the other type are fully rational.

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

Our framework addresses this concern and shows that the extent to which agents coordinate on RE can evolve over time through both adaptive and “eductive” level-k channels. It would be possible to extend our model to include a known fixed proportion of naive agents that follow a known specified rule-of-thumb, or to include a proportion of agents that are fully rational in that they coordinate on RE given the proportions and forecasts of all level-k (and naive) forecasts.⁴

⁴In an earlier version of the paper we included a proportion of agents that were fully

The Cognitive Hierarchy (CH) approach of Camerer, Ho, and Chong (2004), in contrast, allows for a distribution of “k-step” types. In the CH framework this distribution satisfies two assumptions. First, every agent believes, incorrectly, that there are no other agents with equal or higher k-step beliefs; second, every agent knows the exact relative distribution of lower k-step agents. Given these beliefs, k-step agents make optimal decisions conditional on the implied forecasts obtained from those beliefs. Camerer, Ho, and Chong (2004) focus on a family of Poisson distributions that satisfy these assumptions. From our perspective, it is difficult to understand how agents could come to know the distribution of lower k-step types, yet at the same time not realize or consider that there are other agents using equal or higher reasoning steps. These concerns would appear to be even more acute in extensions to repeated or dynamic games.

In the CH approach, as well as in the Reflective Equilibrium approach of García-Schmidt and Woodford (2019) discussed in Section 7, the analysis takes place at a single point in time. In contrast our unified model provides full real-time dynamics for the time paths of level-k forecasts, the proportions of agents using each forecast level, and for the associated time path of the endogenous variables. Our model is suited for analysis of the impact and subsequent dynamics of announced future policy changes, which is often studied in the macroeconomic policy literature.

2 OVERVIEW OF MODEL AND RESULTS

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

The univariate model takes the form $y_t = \gamma + \beta \hat{E}_{t-1} y_t$, where $\hat{E}_{t-1} y_t$ is the average of individual forecasts, made at time $t - 1$, of the variable y_t . Assume $\beta \neq 0, 1$, and for simplicity assume there are no exogenous rational, possessing the knowledge just stated, and we provided the “eductive stability” condition needed for the rational agents to achieve expectational coordination under suitable, strong common knowledge conditions as in Guesnerie (2002).

random shocks so that the REE is $y_t = \bar{y} = (1 - \beta)^{-1} \gamma$. The case $\beta < 0$, with negative expectational feedback, corresponds to the Muth cobweb model of prices in an isolated market with a production lag, while the case $0 < \beta < 1$, with positive expectational feedback, corresponds to a repeated beauty contest or guess-the-average game.

Agents have heterogeneous level- k forecasts. Letting $a_{t-1} = E_{t-1}^0 y_t$ denote the level-0 forecast at $t - 1$, level- k forecasts are defined recursively: $E_{t-1}^k y_t = \gamma + \beta (E_{t-1}^{k-1} y_t)$, for $k = 1, 2, 3 \dots$. Letting ω_{kt} denote the proportion of agents with level- k forecast $E_{t-1}^k y_t$, we have $y_t = \gamma + \beta \sum_{k \geq 0} \omega_{kt} E_{t-1}^k y_t$. Under our unified approach, there are two channels of learning dynamics. First, level-0 dynamics are driven by standard adaptive learning rules updating a_t toward the most recent observation y_t . Second, weights ω_{kt} are updated each period based on replicator dynamic that shifts weight toward the k level providing the most accurate forecast the previous period.⁵

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

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

⁵Our framework could be generalized in various ways. The level- k menu is straightforward to compute, and thus serves to provide natural focal points for forecasts. However, the menu could be extended to include, for example, the average of level- k and level- $k + 1$ forecasts. One could also include level- k calculation costs that are increasing in k . Although extensions like this may prove fruitful in experimental or empirical work, we conjecture that our central findings would be qualitatively unaffected.

and level- k dynamics can work against each other.

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

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

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

We also find evidence for an additional prediction of the unified model:

⁶Readers wanting to focus on the macroeconomic application can omit Section 5 without loss of continuity.

in the wake of announcements subjects may lower their depth of reasoning. We document that some high-level reasoners experience large forecast errors in announcement treatments. This causes a fraction of the high-level reasoners to revise down their depth of reasoning. These downward revisions can make the prevalence of low-level reasoning very persistent

Section 6 extends the unified framework to the New Keynesian model, where we consider monetary policy forward guidance following a persistent stagnation shock, as in Bilbiie (2019). We find that the coupling of adaptive learning and replicator dynamics endogenously induces low level reasoning, substantially reducing the power of monetary policy promises. Section 7 discusses related literature and Section 8 concludes.

3 THE MODEL

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

3.1 THE MODEL

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

The forecasts made by agents with sophistication level k are given by

$E_{t-1}^k y_t$. Aggregate y_t is determined as

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

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. 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 $y_t = \gamma + \beta E_{t-1} y_t$.

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

Agents with higher-order beliefs are assumed to have full knowledge of the model. We recursively define *level- k beliefs* as beliefs that would be optimal if all other agents used level $k - 1$: $E_{t-1}^1 y_t = T(a) \equiv \gamma + \beta a$ and $E_{t-1}^k y_t = T^k(a) \equiv T(T^{k-1}(a))$ for $k \geq 2$. Note that for $k \geq 1$ agents are assumed to know β and γ .⁷

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. $\mathcal{T}(\bar{y}) = \bar{y}$.

⁷This 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 $k \geq 1$.

3.2 ADAPTIVE DYNAMICS

We define *adaptive dynamics* as corresponding to adaptive learning with fixed level- k weights.⁸ 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

$$y_t = \gamma + \beta \sum_{k \geq 0} \omega_k E_{t-1}^k y_t \quad \text{and} \quad a_t = a_{t-1} + \phi(y_t - a_{t-1}), \quad (3)$$

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 *gain* parameter ϕ 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 $i \geq 1$.⁹ Backward-looking rules like (3), as well as anchor and adjustment and trend-following rules, are frequently found to well-describe behavior of laboratory participants in LtFEs as in Hommes (2013). We focus on the specification (3) in order to emphasize the novel features of our framework.

3.3 REPLICATOR DYNAMICS

We next allow agents to revise depth of reasoning over time based on their past forecast performance. Nagel (1995) and Duffy and Nagel (1997) document sluggish updating of reasoning depth in repeated guess-the-average experiments.¹⁰ To capture this sort of updating behavior, we consider the possibility that agents are relatively inattentive to revising their depth of reasoning. We assume a small proportion of those agents using sub-optimal reasoning levels revise their forecast methods, with the proportion 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 (using the same forecast method as in the previous

⁸We 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.

⁹AL can be easily extended to models with observable exogenous shocks, and can allow for heterogeneous forecasts across agents.

¹⁰Khaw, Stevens, and Woodford (2017) also document sluggish adjustment in a different experimental environment.

period), while larger errors incline more agents to “think slow” (revisit and revise their reasoning depth).

We formalize this process by using a replicator dynamic based on 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. Our dynamic shifts weight from suboptimal predictors towards the optimal predictor according to a “rate” function that depends on the forecast error. We define the time t optimal predictor as $\hat{k}(y_t) = \min \arg \min_{k \in \mathbb{N}} |E_{t-1}^k y_t - y_t|$, where the left-most “min” is used to break ties.¹¹

The rate function $r : [0, \infty) \rightarrow [\delta, 1)$ with $\delta \geq 0$ satisfies $r' > 0$.¹² Finally, let ω_{kt} be the weight of level- k beliefs in period t . The system under replicator dynamics is given by

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

$$\omega_{it+1} = \begin{cases} \omega_{it} + \sum_{j \neq \hat{k}(y_t)} r(|E_{t-1}^j y_t - y_t|) \omega_{jt} & \text{if } i = \hat{k}(y_t) \\ (1 - r(|E_{t-1}^i y_t - y_t|)) \omega_{it} & \text{else} \end{cases} \quad (4)$$

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

3.4 UNIFIED DYNAMICS

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

¹¹ $\hat{k}(y_t)$ exists: $k \rightarrow \infty$, $|\beta| < 1$ ($|\beta| > 1$) implies $E_{t-1}^k y_t \rightarrow 0$ ($E_{t-1}^k y_t \rightarrow \infty$).

¹²An 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.

unified dynamics is given as

$$\begin{aligned}
 y_t &= \gamma + \beta \sum_{k \geq 0} \omega_{kt} E_{t-1}^k y_t = \gamma + \beta \sum_{k \geq 0} \omega_{kt} T^k(a_{t-1}) \\
 \omega_{it+1} &= \begin{cases} \omega_{it} + \delta_r \sum_{j \neq \hat{k}(y_t)} r (|T^j(a_{t-1}) - y_t|) \omega_{jt} & \text{if } i = \hat{k}(y_t) \\ (1 - \delta_r r (|T^{\hat{k}(y_t)}(a_{t-1}) - y_t|)) \omega_{it} & \text{else} \end{cases} \quad (5) \\
 a_t &= a_{t-1} + \phi(y_t - a_{t-1}),
 \end{aligned}$$

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

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 $a \neq 0$.

4 PROPERTIES OF THE UNIFIED MODEL

We establish the analytic properties of the unified model, and then turn to simulations for additional insights. These insights are aided by partial analytic results concerning the dependence of \hat{k} on β .

4.1 STABILITY RESULTS

Our central result concerns the stability of the unified model.¹³

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

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

If $\beta < -1$ then odd levels of reasoning introduce negative feedback while even levels result in positive feedback. These countervailing tendencies can result in complex outcomes but also make $\beta < -1$ difficult to analyze.

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

¹³Proofs of all theorems and propositions are found in Appendix A1.

(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 some additional notation is needed. Let $\Delta^n = \{x \in \mathbb{R}^{n+1} : x_i \geq 0 \text{ and } \sum_i x_i = 1\}$ be the n -simplex. 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 . 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 dynamics.

Theorem 3 (Stability of adaptive dynamics). *Let $\delta_r = 0$ 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*
 - (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.*

Items one and two of this theorem are analogous to the results obtained in Theorems 1 and 2; here we also can draw conclusions when $\beta < -1$. The surjectivity of ψ results from the expanding magnitudes and oscillating signs of the β^n . Adaptive dynamics may be written $a_t = \text{constant} + (1 - \phi(1 - \psi))a_{t-1}$, so that the surjectivity of ψ implies that stability and instability may obtain for any value of ϕ . From item 3(b), two additional conclusions can be 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.

4.2 SOME RESULTS ON \hat{k}

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

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

1. *If $|\beta| < 1$ then $K \rightarrow \infty \implies \hat{k} \rightarrow \infty$ and $\hat{k}/K \rightarrow 0$.*
2. *For given K , (a) $\beta \rightarrow -1^- \implies \hat{k} \rightarrow \begin{cases} 1 & \text{if } K \text{ is even} \\ 0 & \text{if } K \text{ is odd} \end{cases}$*
(b) $\beta \rightarrow -1^+ \implies \hat{k} \rightarrow \infty$.

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

4.3 SIMULATIONS OF THE UNIFIED MODEL WITH ANNOUNCEMENTS

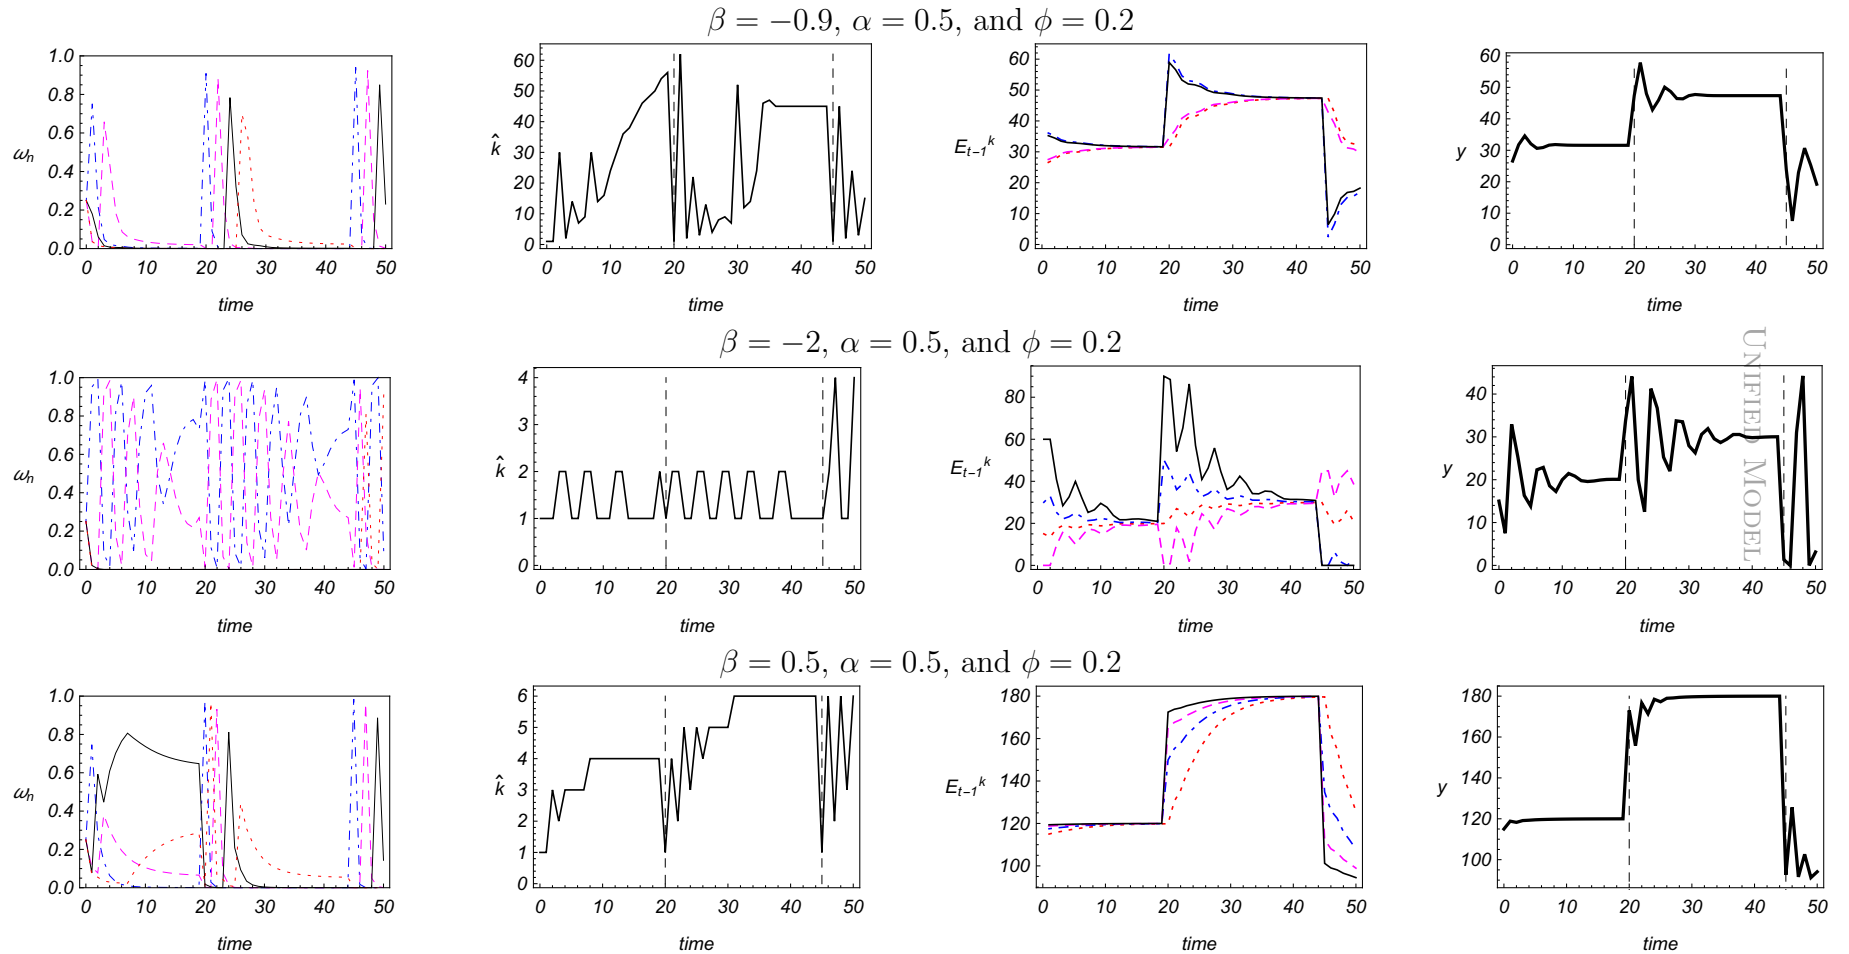
A novel feature of the unified model is that boundedly rational agents can respond to anticipated events by incorporating information about changes in the economic environment. To illustrate, we simulate an economy with a non-zero REE, $\bar{y} > 0$ and a non-negativity constraint on y and $E_{t-1}^k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

Simulations last 50 periods and γ undergoes two announced change: $\gamma = 60$ 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. The announcements are spaced so 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 1 shows the simulated results for the unified model for three different β 's corresponding to the regions of interest identified by our stability theorems. The parameter choices, announcements, and simulation length mirror our experiment. Rows of the figure correspond to different feedback settings. The first plot in each row shows the proportion of agents using the level-0, 1, 2, and 3 predictors; the second plot shows the optimal predictor in use in each period; the third plot shows the level-0, 1, 2, and 3 predictions each period; the fourth plot shows equilibrium dynamics of y_t .

Starting with $\beta = 0.5$, we note three features of the unified dynamics. First, although $y_t = \bar{y}$ for many periods prior to the announcements, there is not instantaneous convergence to the new REE. The existence of low-level reasoners implies that the optimal depth of reasoning in the announcement period is also relatively low (second panel, bottom row). This leads to large forecast errors for those using higher depths of reasoning. Second, in response to these large forecast errors, some high-level reasoners will revise their beliefs *down* to lower levels of reasoning (see the first and second panels, bottom row), leading to another transition period. And third, although agents revise down their depth of reasoning, the proportion who are using a high depth of reasoning remains greater than in the initial periods because not all agents revise their forecasting strategy each period (see first panel, bottom row; the proportion using $k > 3$ is not shown).

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



16

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

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

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

5 THE EXPERIMENT

The unified model makes distinctive predictions for individual expectations and market dynamics. We test these predictions using a standard LtFE experiment. The experiment mirrors the simulated environment of Section 4.3 by having subjects participate in a repeated market for 50 periods, or *rounds*. Subjects forecast the price of a good and are compensated for forecast accuracy. Market price is $p_t = \gamma + \beta \hat{E}_{t-1} p_t + \epsilon_t$, where $\hat{E}_{t-1} p_t$ is the average price forecast across participants and ϵ_t is a small white noise shock, as is standard in LtFE experiments. The shock sequence is the same in all markets and treatments. We adopt a 3×3 experimental design with treatment variables *expectational feedback* ($T\#$) and *timing and size of an announcements* ($A\#$). Treatments are given in Table 1.

Table 1: Experimental Treatments

Feedback Treatments	Announcements Treatments
T1: $\beta = -0.9$	A1: $\gamma = 60$ for $t = 1, \dots, 49$ 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 $t = 45, \dots, 50$.

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

Hypothesis 1 (Stability): *Treatments $\beta < -1$ may not converge or may result in slower rates of convergence compared to treatments $|\beta| < 1$.*

When $|\beta| < 1$, Theorems 1 - 3 imply asymptotic stability of the REE for any specification of the unified dynamic. In addition, the simulations suggest rapid and possibly oscillatory convergence in T1 treatments and monotonic convergence in T3 treatments. In T2 treatments, where $\beta = -2$, results from Theorem 3 and from simulations suggest that asymptotic coordination on the REE is challenging under unified dynamics.¹⁴

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

The announcement treatments allow us to 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. Figure 1 illustrates k-level heterogeneity of individual forecasts; consequently, forecasts should diverge from each other after the announcement. We do not impose, or inform subjects of, a level-0 forecast so coordination on a shared adaptive level-0 forecast is an integral part of the hypothesis.

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

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

¹⁴We remark that with a finite number N of participants, the educative stability condition is relaxed to $-N/(N-1) < \beta < 1$: see Gaballo (2013). The condition here is $-6/5 < \beta < 1$. The A1 treatments are a replication exercise for Bao and Duffy (2016).

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

This hypothesis derives its intuition from results on replicator dynamics – see Theorem 2. In particular, if $|\beta| < 1$ (i.e. T1 and T3) then $\hat{k} \rightarrow \infty$, whereas if $\beta < -1$ (i.e. T2) then \hat{k} is bounded. Finally, we note that the four hypotheses, if true, provide evidence against simpler alternative models. Standard heuristic switching models, for example, are ruled out by Hypothesis 2. Fixed level-k models are ruled out by Hypotheses 3 and 4. Purely adaptive dynamics is ruled out by all hypotheses. Confirmation of the hypotheses supports the unified model over the nested alternatives.

5.1 EXPERIMENT DESCRIPTION

The experiment used a computer based market programmed in oTree.¹⁵ Participants were randomly assigned to groups of six subjects to form markets. Participants acted as expert advisers to firms that produce widgets, and were provided a tutorial on the market environment that included the numerical demand and supply equations. Participants were informed that the price depends on the average expected price across advisers and that prices are subject to small white noise shocks.

Participants were given different stories about the market environment in the positive (T3) and the negative (T1 and T2) feedback cases. The type of feedback in the market is explained in detail with the paper instructions given to participants containing a version of following text: *“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.¹⁶

The payoff for the participant’s predictions is $0.50 - 0.03(p_t - E_{t-1}p_t)^2$ 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.¹⁷

¹⁵See Chen, Schonger, and Wickens (2016) for documentation on oTree. Student recruitment done through ORSEE: Greiner (2015). Figure A9 (Appendix) shows a screenshot of the graphic user interface (GUI). The GUI includes market information in a table, and time series plots of the price and the participants’ previous predictions.

¹⁶We checked for comprehension of the market environment: see Appendix.

¹⁷Negative quantities receive zero cents: see Appendix for further details.

Announcements of changes in γ were made using a pop-up box. The 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. Participants played 50 rounds without time limit. Afterwards, participants were asked which strategy they used, which strategy they believed others used, and which information they found most useful.

5.2 EXPERIMENTAL RESULTS

In total, 372 individuals participated in 62 experimental markets: see Appendix for summary statistics. The first two columns of Figure 2 provide an overview of results from the $T1 \times A3$, $T2 \times A3$, and $T3 \times A3$ treatments. The third column shows the unified model’s fit to the aggregate experimental data. The last column shows the best fit of the nested models within the unified dynamics when either AL, the replicator, or level-k deductions is omitted. The last two columns are discussed in Section 5.2.5, and shown here to demonstrate the necessity of all elements of the unified dynamics.

5.2.1 CONVERGENCE RESULTS The second column of Figure 2 shows convergence properties found across treatments. T1 and T3 quickly converge a few periods after the experiment begins. Markets destabilize following announcements, but quickly re-converge within a few periods. T2 is more volatile: convergence takes longer and individual forecasts continue to vary widely even after near convergence. We measure convergence using three metrics. We discuss the *ratio metric* here and refer to the Appendix for metrics based on mean price discrepancy and mean earnings.

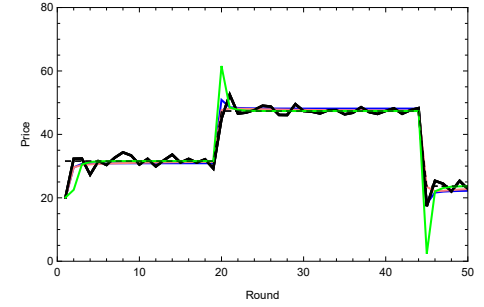

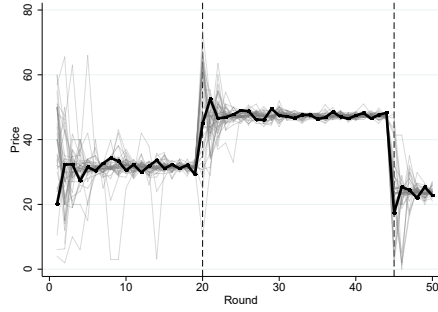
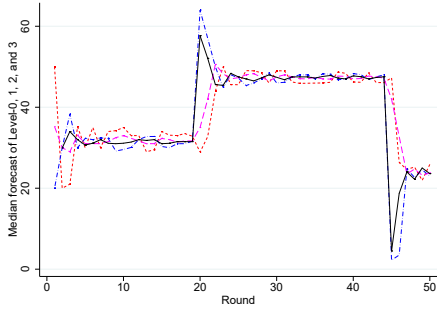
Define a round to be *converged* when price is within ± 3 of the steady-state. The ratio metric applied to consecutive rounds is the proportion of converged rounds. A collection of consecutive rounds has converged if the ratio metric is at least 0.85. By this metric, none of the feedback treatments show convergence within the first five periods or within five periods after the first announcement. However, convergence is achieved for T1 and T3 over rounds 6 to 10, rounds 26 to 30, and overall for the full intervals. For T2, the 85% threshold is never reached. These convergence results, along with other metrics provided in the Appendix, support Hypothesis #1.

Figure 2: Comparing the unified model to experimental data





Experimental Data

Best Fit Models


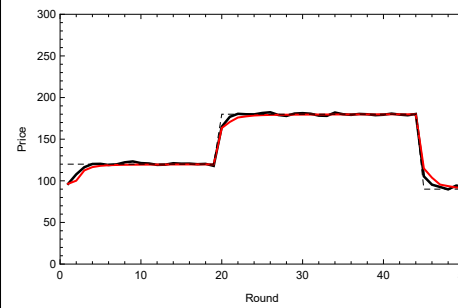
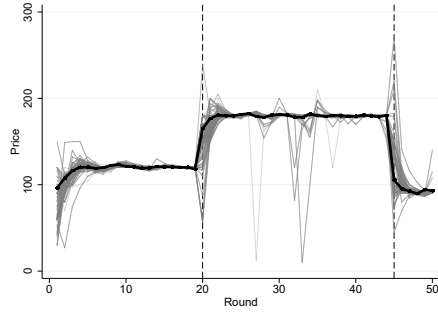

T1 \times A3 ($\beta = -0.9$)



T2 \times A3 ($\beta = -2$)



T3 \times A3 ($\beta = 0.5$)



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

5.2.2 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 provide 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. 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.¹⁸ Using this level-0 forecast for each market, we then construct the implied level-1, 2, 3, and the REE forecasts. Then, we calculate the absolute difference between a subject's forecasts in each round and each of the model implied forecasts. We classify each forecast as either level-1, 2, 3, or the REE according to which has the smallest observed difference. Conflicts in classification, if they arise, are resolved by assigning to the lowest level of reasoning. For the first round, when there is no past history of prices, we use the price from the example on the instructions for the T1 and T2 treatments as the level-0 forecast. The modal forecast given by participants in these treatments is close to this value despite no theoretical reason for why people should choose it. For the T3 treatment, we choose the modal forecast observed in the experimental data in round one as the level-0 forecast.

We stop our classification of types at level-3 deductions because higher levels of deduction become hard to distinguish from the REE forecast in

¹⁸The results are robust to reasonable changes in the definition of level-0 forecast. In the Appendix we reproduce all of our results under the five alternative level-0 assumptions including three constant gain learning specifications and find qualitatively similar results. We also explore one market in detail in the Appendix, which illustrates further how the classification works in practice.

UNIFIED MODEL

the T1 and T3 treatments, and from one another in the T2 treatments in certain settings. We find that approximately 40% of subjects' forecasts that we classified as the REE forecast in a round submit exactly the REE forecast. The remainder are within the ± 3 of it. Therefore, the REE forecast designation likely includes some higher levels of deductions as well.

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

Table 2: Classifying participants' 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)
N	372	348	156	Level-2 x Ann	59.58 (8.806)	12.85 (8.390)
Hypothesis tests of deliberation time regressions				Level-3 x Ann	62.83 (11.80)	22.31 (8.392)
$H_0 : \text{Level-0} - \text{Level-3} = 0$				Cons	39.5 (0.457)	112.68 (4.206)
$H_0 : (\text{Level-0} \times \text{Ann}) - (\text{Level-3} \times \text{Ann}) = 0$				Individual FE	yes	yes
				Round FE	no	yes
				R-squared	0.027	0.253
				N	18,367	18,367

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

The right side of Table 2 provides a logical check on our classifications. It is natural to think that higher levels of deduction require greater cognitive

resources: a person who makes a level-0 forecast might not spend as much time formulating a forecast as someone who makes a level-3 forecast. If our classifications are identifying people who are making level- k deductions we should find a correlation between the time spent deliberating and the depth of reasoning that we identify.

To investigate this, we estimate the following regression model:

$$d_{i,r} = \alpha_i + \omega_r + \sum_k \beta_k I(k)_{i,r} + \sum_k \gamma_k (I(k)_{i,r} \times I(Ann)_r) + \epsilon_{i,r}, \quad (6)$$

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

The regression results confirm our hypothesis. We find that those whom we identify as level-0 spend the least amount of time deliberating on their forecast overall, and in announcement rounds. Those identified as level-3 spend the most amount of time among the classified types in all rounds, and in announcement rounds, with the difference between deliberation times of level-0 and level-3 participants statistically different at standard significance levels. Figure A11, in the Appendix, shows histograms of individual forecasts in round one and in each announcement round for each feedback treatment. For these announcement rounds, Figure A11 and Table 2 show a majority of participants playing level- k or the high level- k /REE forecasts, providing support for Hypothesis #2.

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

Finally, García-Schmidt and Woodford (2019) and Angeletos and Sastry (2021) put forward models of bounded rationality that modify level- k reasoning to rule out oscillating deductions when there is negative expectational feedback. Angeletos and Sastry (2021) writes, “We are not aware of any experimental evidence of this oscillatory pattern. We suspect that it is an unintended “bug” of a solution concept.” In the Appendix, we provide evidence of clear oscillating deductions consistent with classic level- k reasoning for individual participants over time.

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

To test these features of the replicator dynamic, we make use of the announcements in the A2 and A3 treatments, which allow us to identify level- k deductions. Structural change leads to large forecast errors for many participants, and provides counterfactual level- k predictions that can be used to identify revisions to depth of reasoning in the following period.¹⁹

We find evidence consistent with our replicator assumption for all three key aspects: (i) a proportion of subjects do not update their strategy following the announcement period; (ii) subjects who changed strategy experienced larger forecast errors and spent more time deliberating; (iii) a significant proportion of those who do change strategy choose the previous period’s optimal k -level strategy. See Table A4 in the Appendix for details. These findings support Hypothesis #3.

5.2.4 LEVEL-K DYNAMICS The unified model predicts increasing depths of reasoning in T1 and T3 treatments but not in T2 treatments (when convergence is obtained). We show in the Appendix that the distribution of level- k forecasts chosen in treatments with two announcements shifts to the right over time when $|\beta| < 1$ but does not do so when $\beta < -1$. In the latter case, we observe a bifurcation where more level-0 and REE forecasts

¹⁹The classification of forecasts is restricted to level-0, 1, 2, 3, and REE, and is given by the level- k strategy nearest in mean squared error to the submitted forecast.

are played. Overall, these results are consistent with Hypothesis #4.

5.2.5 QUANTITATIVE EVALUATION In this section, we use aggregate price data to compare the fit of the unified model to the fit of simpler alternatives models: a fixed level-k model, a replicator-only model, a pure adaptive learning model, and REE.²⁰

We use our classification of level-k types in period one to initialize the models in each market. The fixed level-k model allows for the level-0 forecast to evolve over time with the proportion of agents using different level-k types fixed to the initial values. The replicator-only model assumes a fixed level-0 forecast but allows for the choice of level-k forecasts to vary over time. The adaptive learning model shuts down the replicator and assumes all agents use the same level zero forecast, which evolves over time as new data become available.

For each version of the model and for each individual market, we compute the mean-squared error (MSE) measured as the average over time of the squared difference between the market price and the price obtained by simulating the model using the fitted values of the associated learning parameters. Table 3 shows the average of the MSEs across the T1×A3, T2×A3, and T3×A3 markets. For given treatments, the average prices across markets, and simulations from each fitted model, are shown in the right two columns of Figure 2. The individual market outcomes are reported in the Appendix.

Using the Wilcoxon ranked-sign test, we compare the individual-market MSEs of the unified model to those of the alternative models. For each of the expectational feedback treatments, the median MSE of the unified model is lower than of the adaptive learning model at the 10% level; and for the T1 and T3 treatments it is lower than for the REE at the 5% level.²¹ The unified model also outperforms the fixed level-k model and the replicator-only model across all three treatments at least the 5% level.

²⁰For each model (except REE), we compute the forecast parameters minimizing the squared error between the simulated data and the experimental data.

²¹For the T2 treatments, the test fails to reject the null hypothesis for equality compared to the REE. We note that the realized market prices are not induced by participants having rational forecasts: see second column of Figure 2.

Table 3: MSE of competing models

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

Notes: Average mean square error (MSE) of five simulated models of aggregate price dynamics compared to experimental market price data. “Rel. RE” reports the MSE of the model relative to RE MSE, i.e., Model MSE/RE MSE. Individual market MSEs that underlie the averages in this table are shown in Table A5 in the Appendix. Models are fit by doing a grid search over values $\alpha \in [0, 2]$ and $\phi \in [0, 1]$.

5.2.6 DISCUSSION The experimental results on individual behavior provide strong support for the four hypotheses. The quantitative evaluation of aggregate data provides evidence that each of underlying mechanisms – adaptive learning, level-k and replicator dynamics – are needed to explain the aggregate data.

6 UNIFIED DYNAMICS IN THE NEW KEYNESIAN MODEL

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

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

We proceed in three steps. First, we introduce the model and the for-

ward guidance policy problem confronting monetary policymakers. Second, we show how a special case of this environment reduces to a univariate model closely related to the model studied in the previous sections. Finally, applying our unified framework to the bivariate NK-model, we examine the implications of an exogenous shock that puts the economy at the zero lower bound. We find that low level-k reasoning can be an endogenous outcome, substantially lowering the power of monetary policy promises.

6.1 FORWARD GUIDANCE POLICY PROBLEM

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

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^n) \quad (7)$$

$$\pi_t = \xi E_t \pi_{t+1} + \kappa x_t, \quad (8)$$

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

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

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

$$\min E_0 \left\{ \frac{1}{2} \sum_{t=0}^{\infty} \xi^t (\pi_t^2 + \psi_x x_t^2) \right\} \quad (9)$$

subject to (7), (8), and $i_t \geq 0$.

Optimal discretionary policy is to set the nominal rate at zero in the stagnation state. To study forward guidance, we follow Bilbiie (2019) and assume that policymakers engage in a partial commitment strategy: in

state S the bank announces that it will continue to hold the interest rate at zero beyond the end of state S , i.e., provide forward guidance. The implementation of this policy involves a probabilistic return to normalcy: after the natural rate returns to its normal value of r_N policy makers continue to hold the interest rate at zero each period with probability $1 - \nu$.

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

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

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

6.2 UNIFIED DYNAMICS IN THE NEW KEYNESIAN MODEL

To develop unified dynamics in this multivariate environment, we, as usual, interpret $E_t x_{t+1}$ and $E_t \pi_{t+1}$ in (7) and (8) as the average expectations across agents of output gap and inflation in period $t + 1$. In this way, equations (7) and (8) can be taken as the current period best response functions over which the agents do level-k deductions.²² This is the standard assumption employed in macroeconomic laboratory experiments that test expectation formation in the New Keynesian model such as in Mokhtarzadeh and Petersen (2021) and Kryvtsov and Petersen (2021).²³

To illustrate how level-k deductions work in this setting, and to facilitate connections to our previous analysis, we begin by deriving these deductions in the special case $\kappa = 0$ in which prices are fixed, so that inflation and inflation expectations are zero, and no forward guidance, i.e. $\nu = 1$. The general case is derived in the Appendix.

Assume that in period $t \geq 0$ the economy is (still) in state S and discretionary monetary policy is pursued. We assume that, in state S ,

²²In environments with long-lived agents there are a variety of ways to model decision-making: for an alternative implementation see Preston (2005), and for further discussion see Evans and McGough (2020) and Evans and McGough (2021).

²³Angeletos and Lian (2018) and Farhi and Werning (2019) develop alternate implementations of k-level deductions in New Keynesian environments.

level-0 agents hold forecasts $E_t^0[x_{t+1}|S] = a_{t-1}^x$, and that in state N , level-0 agents hold forecasts $E_t^0[x_{t+1}|N] = 0$.²⁴ We continue to assume that level-1 agents assume that all agents are level-0. Thus for given level zero expectations a_{t-1}^x of the output gap in state S , level-1 forecasts are obtained as follows:

$$\begin{aligned} E_t^1[x_{t+1}|S] &= (1 - \delta)E_{t+1}^0[x_{t+2}|S] + \delta \underbrace{E_{t+1}^0[x_{t+2}|N]}_0 + \sigma(1 - \delta)r_S \\ &= (1 - \delta)a_{t-1}^x + \sigma(1 - \delta)r_S. \end{aligned}$$

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

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

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

Continuing in this way, we can define $E_t^0[x_{t+1}|S] = a_{t-1}^x$,

$$\begin{aligned} E_t^1[x_{t+1}|S] &= T(a_{t-1}^x|S) \equiv (1 - \delta)\sigma r_S + (1 - \delta)a_{t-1}^x, \\ E_t^k[x_{t+1}|S] &= T^k(a_{t-1}^x|S) \equiv T(T^{k-1}(a_{t-1}^x|S)) \text{ for } k \geq 2. \end{aligned}$$

Combining these definitions with the IS equation (7) and substituting in aggregate beliefs yields the realized value of x in state S as a function of level-0 beliefs, i.e. $x_t = \mathcal{T}(a_{t-1}^x|S)$, where

$$\mathcal{T}(a^x|S) = \sigma \cdot r_S \left(1 + \delta^{-1}(1 - \delta) \sum_{k \geq 0} \left(1 - (1 - \delta)^k \right) \omega_k \right) + \left(\sum_{k \geq 0} (1 - \delta)^k \omega_k \right) a^x. \quad (10)$$

Note \mathcal{T} is linear in a^x and (10) is the analog to equation (2) in Section 2.

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

²⁴The narrative here is that there has been an extended period in which the economy has been in state N , and agents have learned that in state N output gap is zero. By contrast, state S (or F) is novel/unusual, and agents have to learn about it.

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

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

$$\mathbb{L}_t^k(z_t) = \left| \hat{E}_{t-1}^k[\pi_t|z_t] - \pi_t \right| + \psi_x \left| \hat{E}_{t-1}^k[x_t|z_t] - x_t \right|, \quad (11)$$

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

$$\hat{k}(x_t, \pi_t, z_t) = \min \arg \min_{k \in \mathbb{N}} \mathbb{L}_t^k(z_t), \quad (12)$$

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

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

$$\begin{aligned} x_t &= \sum_{k \geq 0} \omega_{kt} E_t^k x_{t+1} - \sigma(i_t - \sum_{k \geq 0} \omega_{kt} E_t^k \pi_{t+1} - r_t^n) \\ \pi_t &= \xi \sum_{k \geq 0} \omega_{kt} E_t^k \pi_{t+1} + \kappa x_t \\ \omega_{it+1} &= \begin{cases} \omega_{it} + \sum_{j \neq \hat{k}(y_t)} r(\mathbb{L}_t^j(z_t)) \omega_{jt} & \text{if } i = \hat{k}(x_t, \pi_t, z_t) \\ (1 - r(\mathbb{L}_t^k(z_t))) \omega_{it} & \text{else} \end{cases} \\ a_t^x &= a_{t-1}^x + \phi(x_t - a_{t-1}^x) \text{ and } a_t^\pi = a_{t-1}^\pi + \phi(\pi_t - a_{t-1}^\pi), \end{aligned} \quad (13)$$

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

6.3 LOW-LEVELS OF DEDUCTION AND FORWARD GUIDANCE

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

$$E_t^k x_{t+1} = \nu^{-1} (1 - (1 - \nu)^k) \sigma r_N + (1 - \nu)^k a^x. \quad (14)$$

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


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

In the general case, with $\kappa > 0$, low levels of reasoning are persistent in the face of adverse shocks coupled with forward guidance. This is the result of competing tensions. For our calibrations, a negative shock coupled with forward guidance results in forecasts that are non-monotonic in k . Consider inflation forecasts $E_0^k \pi_1$ at the time of the shock and for different values of k . At low level- k , both the shock and policy have small contemporaneous effects because agents do not fully consider the potential persistence of the shock nor future policy changes. Increasing k at first generates pessimism as the negative impact of the persistent shock dominates the positive impact of future policy. For sufficiently high k future policy is more salient and can dominate the pessimism. Thus some low level- k forecasts and some high level- k forecasts generate similar predictions, which removes the incentive

under the replicator dynamics for low-level reasoners to revise upward.

Suppose the economy the economy unexpectedly enters state S at the beginning of period $t = 0$. Figure 3 illustrates the attendant non-monotonicity in (reasoning level) k of forecasts of period 1 inflation. We use the Eggertsson and Woodford (2003) calibration: $\xi = 0.99$, $\kappa = 0.2$, $\sigma = 0.5$, $\delta = 0.1$, and $r_N = 0.01$, with a large shock of $r_S = -0.01$. The level-0 forecast is $a^x = a^\pi = 0$. The three solid lines are the same in each panel and show the level- k inflation forecasts for the three different forward guidance promises ν . The forward guidance promise ν^* corresponds to optimal promise under RE when ψ_x equals the welfare theoretic value of 0.00254.²⁵ The solid lines clearly illustrate the non-monotonicity (the same is true for the output gap forecast). Level-500 is approximately RE.

Figure 3: Non-monotonicity of level- k forecasts



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

The four panels in Figure 3 differ in the assumed distribution of level- k

²⁵The welfare theoretic value is κ/θ , where θ is the elasticity of demand with respect to price faced by the monopolistically competitive firms in the economy.

types which generate different values of inflation in period $t = 1$. These values are indicated by the dashed lines. For each ν , the level- k forecast with lowest absolute error is indicated by the large dot. The top panels show the results for uniform distributions of level- k types: $[0, 499]$ (on left) and $[0, 3]$ (on right). The left bottom panel shows the results for the distribution of types observed in our T3×A2 and A3 experiments for the first announcement with proportions of level-0, 1, 2, 3, and REE forecasts of 25.4%, 56.8%, 4.2%, 1.7%, and 11.9%, respectively. In the last panel a weight $1/2$ is placed on level-0 and a weight of $1/2$ is placed on level-500.

In all four cases the optimal level- k is smaller than the largest k in use, and for some larger forward guidance promises (smaller ν) there are double crossing of level- k forecasts and realized inflation. This shows that low and high level- k forecasts can produce similarly small forecast errors. Thus low level- k strategies do not lose many users even when the optimal k is higher.

Figure 4 compares inflation and output dynamics under RE (dashed black) and under unified dynamics (solid black), using the same the initial distribution of level- k types as the bottom left panel of Figure 3. We set the gain to $\phi = 0.2$ and the replicator parameter to $\alpha = 500$: this corresponds to a loss of nearly 90% of the users of a level- k strategy in a single period if the absolute forecast error for inflation is 1%. By period 10, few use the poorly-performing RE forecast; however, because low-level- k forecasts remain good, the use of level-0, 1, 2, and 3 forecasts is still above 30%.


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

6.4 DISCUSSION

We conclude with observations made by Farhi and Werning (2019), who write,

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

Figure 4: Forward guidance under unified dynamics simulation



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

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

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

7 RELATED LITERATURE

In addition to the level- k literature and the cognitive hierarchy approaches discussed in the Introduction, the unified model draws on several literatures. The eductive approach introduced in Guesnerie (1992) examines the inherent difficulty for rational agents, who fully know the structural model, to coordinate on REE. This coordination requires extremely strong common knowledge assumptions, not only of the structure but also of the rationality of other agents; even then coordination on an REE is only possible if the structure satisfies certain “eductive stability” conditions. These are closely related to the iterative expectational stability conditions de-

veloped in Evans (1985), a connection developed explicitly in Evans and Guesnerie (1993). Our level- k forecasts are obtained analogously, starting from level-0 forecasts, using iterations based on the structure.

Our level-0 learning is based on the adaptive learning literature developed in Bray and Savin (1986), Marcet and Sargent (1989), Evans (1989) and Evans and Honkapohja (2001). AL is a versatile technique that has been applied in both nonexperimental and experimental settings. For a wide range of models, AL can converge over time to REE.²⁶ Because AL does not require knowledge of the model’s structural parameters, it provides a natural level-0 benchmark for deriving level- k forecasts.²⁷

The behavioral heterogeneous expectations literature, e.g. Brock and Hommes (1997), De Grauwe (2012) and Hommes (2013), considers ex ante homogeneous agents selecting from a menu of forecast rules, resulting in ex post heterogeneity of forecasts. In our setting the menu is the full set of level- k forecasts, with agents’ choices based on recent forecast performance.

Our model shares elements with the Reflective Equilibrium notion proposed by García-Schmidt and Woodford (2019), which features a continuous version of level- k forecasts parameterized by a finite “degree of reflection” n , viewed as the mean level of reasoning. RE is obtained as $n \rightarrow \infty$, but they argue a Reflective Equilibrium with a finite degree of reflection is more realistic. As they note, this approach is similar to the “calculation equilibrium” analyzed in Evans and Ramey (1992), in which agents revise expectations of future paths recursively, with increased calculation costs from higher levels of recursion balanced against reduced forecast errors. For simplicity we do not include these costs.

Our LtFE shares important elements with the laboratory experiments of Fehr and Tyran (2008), Heemeijer, Hommes, Sonnemans, and Tuinstra (2009), and Bao, Hommes, Sonnemans, and Tuinstra (2012). Each study tests for convergence to an REE in an LtFE setting. Bao, Hommes, Sonnemans, and Tuinstra (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 detailed structure of

²⁶Sargent (2008) argues REE can be viewed as emergent outcomes from learning.

²⁷Evans, Guesnerie, and McGough (2018) show that though the RBC model is not deductively stable, it is stable under AL.

the model, and level- k forecasts are not studied. Using a pricing game, Heemeijer, Hommes, Sonnemans, and Tuinstra (2009) find that negative feedback engenders stability while positive feedback can lead to endogenous fluctuations.²⁸ Fehr and Tyran (2008) study speed of convergence in a pricing game with different feedback treatments, which they refer to as strategic substitutability and strategic complementarity.

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

In Anufriev, Duffy, and Panchenko (2022), subjects forecast two variables whose realizations are dependent on each other. They argue that both AL and level- k reasoning are necessary to fit their data. By contrast, our unified approach provides sharp predictions about revisions to depth of reasoning and the impact of anticipated events, and our experiment tested these predictions.

8 CONCLUSION

The union of behavioral heterogeneity, adaptive learning, and level- k reasoning brings together three assumptions that enjoy wide experimental support. We show how evolving level- k beliefs naturally fit common forms of bounded rationality studied in macroeconomic environments. One of our key findings 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, including prominent applications to the forward guidance puzzle.

The unified approach has a number of features that we find appealing. It naturally balances adaptive and eductive approaches to expectations formation by providing agents with some structural knowledge while also

²⁸See Sutan and Willinger (2009) for level- k experiments with negative feedback.

assuming they update beliefs over time as new data become available; it is amenable to theoretical analysis and yields both intuitive and surprising results; it is shown to be supported in experiments; and it is easily and naturally adapted to many general economic models.

REFERENCES

- ANGELETOS, G.-M., AND C. LIAN (2018): “Forward guidance without common knowledge,” *American Economic Review*, 108(9), 2477–2512.
- ANGELETOS, G.-M., AND K. A. SASTRY (2021): “Managing Expectations: Instruments Versus Targets,” *The Quarterly Journal of Economics*, 136(4), 2467–2532.
- ANUFRIEV, M., J. DUFFY, AND V. PANCHENKO (2022): “Learning in two-dimensional beauty contest games: Theory and experimental evidence,” *Journal of Economic Theory*, 201, 105417.
- ARIFOVIC, J., S. SCHMITT-GROHÉ, AND M. URIBE (2018): “Learning to live in a liquidity trap,” *Journal of Economic Dynamics and Control*, 89, 120–136.
- BAO, T., AND J. DUFFY (2016): “Adaptive versus educative learning: Theory and evidence,” *European Economic Review*, 83, 64–89.
- BAO, T., C. HOMMES, J. SONNEMANS, AND J. TUINSTRRA (2012): “Individual expectations, limited rationality and aggregate outcomes,” *Journal of Economic Dynamics and Control*, 36(8), 1101–1120.
- BIANCHI, F., M. LETTAU, AND S. C. LUDVIGSON (2022): “Monetary policy and asset valuation,” *The Journal of Finance*, 77(2), 967–1017.
- BIANCHI, F., AND L. MELOSI (2014): “Dormant shocks and fiscal virtue,” *NBER Macroeconomics Annual*, 28(1), 1–46.
- BILBIE, F. O. (2019): “Optimal forward guidance,” *American Economic Journal: Macroeconomics*, 11(4), 310–345.
- BOSCH-DOMENECH, A., J. G. MONTALVO, R. NAGEL, AND A. SATORRA (2002): “One, two,(three), infinity,...: Newspaper and lab beauty-contest experiments,” *American Economic Review*, 92(5), 1687–1701.
- BRANCH, W. A., AND B. MCGOUGH (2008): “Replicator dynamics in a cobweb model with rationally heterogeneous expectations,” *Journal of Economic Behavior & Organization*, 65(2), 224–244.
- BRAY, M. M., AND N. E. SAVIN (1986): “Rational expectations equilibria, learning, and model specification,” *Econometrica*, 54, 1129–1160.
- BROCK, W. A., AND C. H. HOMMES (1997): “A rational route to randomness,” *Econometrica*, pp. 1059–1095.
- CAMERER, C. F., T.-H. HO, AND J.-K. CHONG (2004): “A cognitive hierarchy model of games,” *The Quarterly Journal of Economics*, 119(3), 861–898.
- CHEN, D. L., M. SCHONGER, AND C. WICKENS (2016): “oTree—An open-source platform for laboratory, online, and field experiments,” *Journal of Behavioral and Experimental Finance*, 9, 88–97.
- COSTA-GOMES, M. A., AND V. P. CRAWFORD (2006): “Cognition and behavior in two-person guessing games: An experimental study,” *American Economic Review*, 96(5), 1737–1768.
- DE GRAUWE, P. (2012): *Lectures on behavioral macroeconomics*. Princeton University Press.
- DUFFY, J., AND R. NAGEL (1997): “On the robustness of behaviour in experimental ‘beauty contest’ games,” *Economic Journal*, 107(445), 1684–1700.

- EGGERTSSON, G. B., AND M. WOODFORD (2003): “Zero bound on interest rates and optimal monetary policy,” *Brookings papers on economic activity*, 2003(1), 139–233.
- EVANS, G. W. (1985): “Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models,” *The Quarterly Journal of Economics*, 100, 1217–1233.
- (1989): “The Fragility of Sunspots and Bubbles,” *Journal of Monetary Economics*, 23, 297–317.
- EVANS, G. W., AND R. GUESNERIE (1993): “Rationalizability, strong rationality, and expectational stability,” *Games and Economic Behavior*, 5(4), 632–646.
- EVANS, G. W., R. GUESNERIE, AND B. MCGOUGH (2018): “Eductive Stability in Real Business Cycle Models,” *Economic Journal*, 129(618), 821–852.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and expectations in macroeconomics*. Princeton University Press.
- EVANS, G. W., S. HONKAPOHJA, AND K. MITRA (2009): “Anticipated fiscal policy and adaptive learning,” *Journal of Monetary Economics*, 56(7), 930–953.
- EVANS, G. W., AND B. MCGOUGH (2020): “Adaptive Learning in Macroeconomics,” *Oxford Research Encyclopedia of Economics and Finance*, <https://doi.org/10.1093/acrefore/9780190625979.013.508>.
- (2021): “Agent-level Adaptive Learning,” *Oxford Research Encyclopedia of Economics and Finance*, <https://doi.org/10.1093/acrefore/9780190625979.013.620>.
- EVANS, G. W., AND G. RAMEY (1992): “Expectation Calculation and Macroeconomic Dynamics,” *American Economic Review*, 82, 207–224.
- FARHI, E., AND I. WERNING (2019): “Monetary policy, bounded rationality, and incomplete markets,” *American Economic Review*, 109(11), 3887–3928.
- FEHR, E., AND J.-R. TYRAN (2008): “Limited rationality and strategic interaction: the impact of the strategic environment on nominal inertia,” *Econometrica*, 76(2), 353–394.
- GABALLO, G. (2013): “Eductive learning and the rationalizability of oligopoly games,” *Economics Letters*, 120(3), 401–404.
- GALI, J., AND M. GERTLER (1999): “Inflation dynamics: A structural econometric analysis,” *Journal of Monetary Economics*, 44(2), 195–222.
- GARCÍA-SCHMIDT, M., AND M. WOODFORD (2019): “Are low interest rates deflationary? A paradox of perfect-foresight analysis,” *American Economic Review*, 109(1), 86–120.
- GIBBS, C. G., AND M. KULISH (2017): “Disinflations in a model of imperfectly anchored expectations,” *European Economic Review*, 100, 157–174.
- GOY, G., C. HOMMES, AND K. MAVROMATIS (2020): “Forward guidance and the role of central bank credibility under heterogeneous beliefs,” *Journal of Economic Behavior & Organization*.
- GREINER, B. (2015): “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE,” *Journal of the Economic Science Association*, 1(1), 114–125.
- GUESNERIE, R. (1992): “An Exploration of the Eductive Justifications of the Rational-Expectations Hypothesis,” *American Economic Review*, 82(5), 1254–1278.
- (2002): “Anchoring economic predictions in common knowledge,” *Econometrica*, 70(2), 439–480.
- HEEMELJER, P., C. HOMMES, J. SONNEMANS, AND J. TUINSTRA (2009): “Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation,” *Journal of Economic Dynamics and Control*, 33(5), 1052–1072.
- HO, T.-H., C. CAMERER, AND K. WEIGELT (1998): “Iterated dominance and iterated best response in experimental” p-beauty contests,” *American Economic Review*, 88(4), 947–969.

- HOMMES, C. (2011): “The heterogeneous expectations hypothesis: Some evidence from the lab,” *Journal of Economic Dynamics and Control*, 35(1), 1–24.
- (2013): *Behavioral rationality and heterogeneous expectations in complex economic systems*. Cambridge University Press.
- HOMMES, C., J. SONNEMANS, J. TUINSTRAS, AND H. VAN DE VELDEN (2007): “Learning in cobweb experiments,” *Macroeconomic Dynamics*, 11(S1), 8–33.
- JACKSON, A. L. (2005): “Disinflationary Boom Reversion,” *Macroeconomic Dynamics*, 9(4), 489–515.
- KAHNEMAN, D. (2011): *Thinking fast and slow*. Farrar, Straus and Giroux, New York.
- KHAW, M. W., L. STEVENS, AND M. WOODFORD (2017): “Discrete adjustment to a changing environment: Experimental evidence,” *Journal of Monetary Economics*, 91, 88–103.
- (2019): “Adjustment dynamics during a strategic estimation task,” ”Working paper”.
- KRYVTSOV, O., AND L. PETERSEN (2021): “Central bank communication that works: Lessons from lab experiments,” *Journal of Monetary Economics*, 117, 760–780.
- LUCAS, JR., R. E. (1972): “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, 4, 103–124.
- MARCEY, A., AND T. J. SARGENT (1989): “Convergence of least squares learning mechanisms in self-referential linear stochastic models,” *Journal of Economic theory*, 48(2), 337–368.
- MAUERSBERGER, F., AND R. NAGEL (2018): “Levels of reasoning in Keynesian Beauty Contests: a generative framework,” in *Handbook of computational economics*, vol. 4, pp. 541–634. Elsevier.
- MOKHTARZADEH, F., AND L. PETERSEN (2021): “Coordinating expectations through central bank projections,” *Experimental Economics*, 24, 883–918.
- MUTH, J. F. (1961): “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29, 315–335.
- NAGEL, R. (1995): “Unraveling in guessing games: An experimental study,” *American Economic Review*, 85(5), 1313–1326.
- NAGEL, R., C. BÜHREN, AND B. FRANK (2017): “Inspired and inspiring: Hervé Moulin and the discovery of the beauty contest game,” *Mathematical Social Sciences*, 90, 191–207.
- PRESTON, B. (2005): “Learning About Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1(2), 81–126.
- SARGENT, T. J. (2008): “Evolution and Intelligent Design,” *American Economic Review*, 98, 5–37.
- SETHI, R., AND R. FRANKE (1995): “Behavioural heterogeneity under evolutionary pressure: macroeconomic implications of costly optimisation,” *Economic Journal*, 105(430), 583–600.
- SUTAN, A., AND M. WILLINGER (2009): “Guessing with negative feedback: An experiment,” *Journal of Economic Dynamics and Control*, 33(5), 1123–1133.
- WEIBULL, J. W. (1997): *Evolutionary game theory*. MIT press.
- WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.