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

October 14, 2020

Preliminary

Abstract

We propose and experimentally test a model of boundedly rational decision making that combines adaptive learning, eductive learning, and level-k reasoning. We assume that heterogeneous agents engage in level-k deductions, where level-0 agents are described by adaptive learners and level-infinity agents are consistent with eductive learners. Agents revise their depth reasoning in response to forecast error following a type of replicator dynamic. We show that the unified model can rationalize observed behavior in Learning-to-Forecast Experiments when laboratory participants have significant knowledge of the market environment such as faster or slower convergence to the rational expectations equilibrium than expected if people were purely adaptive learners or were acting rationally. In addition, the model makes predictions about how agents forecast anticipated events. We test these predictions by introducing announced structural changes into an otherwise standard forecasting game in a cobweb market environment that permits both positive (strategic complements) and negative (strategic substitutes) feedback. We find that the model well describes the experimental data across all treatments and rationalizes behaviors observed in other Learning-to-Forecast and beauty contest experiments.

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

Key Words: Expectations; Learning; Stability; Behavioral; Experiments.

*We thank Yuqi Angel Jiang and Laura Flakowski for research assistant support. This paper has benefited from comments received from Valentyn Panchenko and Isabelle Salle and in talks at the University of Amsterdam, Universitat Rovira i Virgili, Australian National University, the University of Adelaide, the 26th Symposium of SNDE at Keio University, and the Expectations and Dynamic Macroeconomic Models Conference hosted by the University of Birmingham. Financial support from National Science Foundation Grant No. SES-1559209, UNSW Business School, and The University of Sydney is gratefully acknowledged.

1 Introduction

The assumption of rational expectations (RE) continues to come under scrutiny in macroeconomics and finance models in which RE plays a central role. RE imposes strong assumptions on agents' knowledge and cognitive abilities that calls into question the plausibility and robustness of some model predictions. Widely used boundedly-rational alternatives such as adaptive learning, as in Bray and Savin (1986), Marcet and Sargent (1989), and Evans and Honkapohja (2001) or behavioral approaches, as in De Grauwe (2012) or Hommes (2013), attempt to overcome these concerns by modeling expectations in disciplined ways that often reflect observed behavior in laboratory experiments (Hommes, Sonnemans, Tuinstra and Van De Velden (2007), Hommes (2011), Bao, Hommes, Sonnemans and Tuinstra (2012)) or surveys of forecasters (Branch (2004)). However, a limitation of several of these alternatives is that they focus solely on non-strategic behavior and provide little guidance on how agents may forecast general equilibrium effects of current or future events.[1] A drawback that has limited their usefulness in tackling recent puzzles in RE modeling such as those related to monetary policy and the zero lower bound.[2]

An alternative to these adaptive frameworks is the eductive approach of Guesnerie (1992) and Guesnerie (2002). In eductive learning, agents are fully rational and fully understand the structure of the economy, and furthermore, this is common knowledge. But agents still must independently coordinate on a unique outcome. The conditions under which coordination can take place rules out a number of possible equilibria that would satisfy RE conditions. When the RE equilibrium (REE) is eductively stable, though, the predictions is for instantaneous coordination on the REE, which strongly conflicts with experimental evidence.

However, the notional way in which coordination occurs under eductive reasoning, through iterated deletion of strictly dominated strategies, does have experimental support. Studies of "beauty contest" or "guess the average" games such as Nagel (1995), Duffy and Nagel (1997), Ho, Camerer and Weigelt (1998), Bosch-Domenech, Montalvo, Nagel and Satorra (2002), and Costa-Gomes and Crawford (2006) find strong support for a type of iterated deductions known as level-k reasoning. In these experiments, participants are observed to make a finite number of deductions using their common knowledge of the game but assuming their opponents are less than rational. Therefore, instantaneous coordination on the REE does not occur because agents do not iterate out indefinitely, but the basic logic of eductive reasoning is on display.

[1] There are some notable exceptions including some adaptive learning-based papers that have proposed ways to model anticipated changes about future policy such as Evans, Honkapohja and Mitra (2009) for anticipated fiscal policy and Evans and McGough (2018) for anticipated monetary policy. However, the forecasts in these cases are partial equilibrium. Recently, Gibbs and Kulish (2017) propose a framework for modeling boundedly-rational general equilibrium forecasts for current and future events but it is not fully dynamic and it is nested by the framework we propose here.

[2] A particularly prominent example is the so-called Forward Guidance Puzzle (Del Negro, Giannoni and Patterson (2012)), where standard RE structural models that are used to evaluate monetary policy produced implausible predictions about the power of forward guidance at the zero lower bound.

Bao and Duffy (2016) explicitly attempts to test whether the insights from eductive reasoning have predictive power in the context of a cobweb model in which there is negative feedback. They leverage the fact that when there is negative feedback the unique REE is always learnable by adaptive agents but coordination by eductive agents is only obtained if the negative feedback is appropriately bounded. They find that the bound for eductive stability is predictive for a change in subject forecasting behavior but that behavior cannot be described as either adaptive or eductive. They argue it is a mixture of both.

In this paper we propose a unified model of bounded rationality that nests adaptive and eductive learning as special cases.[3] We marry these two concepts by way of level-k reasoning following Nagel (1995). We show that within an economic environment that depends on one-step-ahead expectations that we can rationalize a host of behaviors observed in Learning-to-Forecast Experiments (LtFEs) and repeated beauty contest games. The model also provides a boundedly rational theory that makes general equilibrium predictions for anticipated events, which we explicitly test using a LtFE.

Our model unifies four different strands of the macro behavioral literature: behavioral heterogeneity, adaptive learning, level-k reasoning, and eductive reasoning. We assume behavioral heterogeneity following Hommes (2013), where agents are ex ante identical but can choose among different forecasting strategy over time, which leads to cross-sectional heterogeneity in beliefs. In contrast to much of the earlier work in this area, we assume that the number of forecasting strategies agents consider is infinite. The menu of strategies is book-ended by adaptive learning and eductive reasoning, while intermediate strategies follow level-k depths of reasoning. Specifically, the least sophisticated forecasting strategy, level-0, uses a salient value, given the structure of the model, without strategic considerations, e.g. an adaptive process. Level-1 agents choose a forecast optimally given that all other agents are level-0 using their knowledge of the economic environment, and level-k agents, for $k = 2, 3, \ldots$, are defined inductively as agents who choose the best response assuming that other agents are level $k - 1$. In the limit, as k goes to infinity, the terminal forecasting strategy is defined by eductive reasoning.

Each agent's strategy choice is governed by a kind of replicator dynamic. Each period a subset of agents whose strategies are performing poorly undertake a counterfactual exercise where they compare how well all possible strategies would have performed in the most recent period. They then select the best strategy observed from this exercise. The unified model, therefore, features two types of learning: adaptive learning as a possible forecasting strategy and learning over the depth of reasoning, i.e. the 'k' in level-k.

In a cobweb market environment, the unified model makes some sharp predictions for when and how fast convergence to the REE is obtained. Consistent with the results of Bao and Duffy (2016), convergence depends on the expectational feedback in the market. When feedback, given by the parameter β, is bounded between one and negative one, we prove

[3] This allows us to explicitly model the conjecture put forward by Evans (2001) that expectation formation is both an adaptive and an eductive process.

convergence under learning only, under the replicator only, and show by simulation, convergence of the unified model when both mechanisms are in operation. When $\beta < -1$, both convergence and non-convergence are possible depending on the speed of adjustment across both dimensions, which provides an explanation for Bao and Duffy's mixed experimental results. Small changes in how fast agents update beliefs can lead to drastically different dynamics when $\beta < -1$.

In addition, the unified model predicts different speeds of convergence for when feedback is positive compared to when it is negative. When $\beta > 0$ such that the market exhibits strategic complementarity, there is slower convergence compared to when $-1 < \beta < 0$ and the market exhibits strategic substitutability. This is consistent with the experimental studies of Fehr and Tyran (2008) and Bao et al. (2012) but the unified model sheds new light on the possible mechanisms. In particular, Fehr and Tyran argue that what explains these observed differences in convergence speed is the large forecast errors generated in environments of strategic substitutability relative to strategic complementarity. They argue that larger errors, ceteris paribus, in the former case cause learning agents to update beliefs more quickly driving the market to equilibrium. However, under the unified model, there is an additional force that can slow or speeds up convergence, which is the selection of a depth of reasoning through the replicator dynamic. When there is heterogeneity among agent's depths of reasoning, the optimal level of k that agents move towards depends on the expectational feedback in the market. When there is positive feedback, the optimal level-k is low relative to the level-k strategy in use at the time whereas the opposite is true with negative feedback. Lower depths of reasoning forecasts are farther away from the equilibrium and vice versa, which slows or speeds up convergence.

The addition of the strategic consideration in unified model also can reconcile Fehr and Tyran's findings with Bao and Duffy (2016) finding that very slow or non-convergence occurs when there is extreme negative feedback. The errors are even larger in these cases, which given Fehr and Tyran's hypothesis should make convergence happen even faster. But this is offset by the strategic actions that agents take when attempting to coordinate in this environment because of level-k deductions.

We take the model to a laboratory to empirically validate its predictions using a standard Learning-to-Forecast Experiment. The experimental design largely follows that of Bao and Duffy (2016). We study a cobweb market model with a non-storable good produced with a production lag. The market price depends on one-period ahead expectations, which are supplied by experimental participants who have full information of the market structure, which nests the common beauty contest game but in an explicitly dynamic setting. Participant are paid based on the accuracy of their forecasts.

The key treatment in our experiment is the introduction of announced changes to the experimental market that occur at irregular intervals. The announcements are akin to embedding a beauty contest game within the market game allowing us to see how participants incorporate new information into their forecasts. The advantage of this approach is that

the periods leading up to an announcement provide data for participants - and for us as the researchers - to identify plausible level-0 beliefs from which all other level-k forecasts are derived. The unified theory provides clear predictions for the distribution of forecasts observed in announcement periods as well as how people should revise their depth of reasoning in subsequent periods.

We find that the unified model well-captures the experimental data. We find strong evidence for adaptive and level-k type reasoning expectations. In particular, in announcement periods we can classify between 70% to 75% of participants as either level- 0, 1, 2, 3 or as those who guessed the REE depending on how we measure. In addition, we are able to document revisions to the depth of reasoning for and across individuals that correspond with the unified theory's conditions for convergence to the REE. When $-1 < \beta < 1$, we see the distribution of the depth of reasoning shift to the right comparing across the two separate announcements. However, when $\beta < -1$, we find that depth of reasoning bifurcates with some agents jumping to the REE and other reverting to purely adaptive behavior. This bifurcation delays or even prevents convergence to the REE in these markets consistent with the findings of Bao and Duffy (2016).

Finally, we extend the unified model to an environment with an infinite forecasting horizon. We show that level-k deductions here have a natural interpretation of how forward-looking an agent believes others to be. When there are anticipated changes in the market structure at some known date in the future, the contemporaneous response depends on whether the agent believes that the other agents will respond to this news today. For example, if everyone is level-1, then the individual belief is that everyone else is level-0 and will not respond to new information about a future event. The best response to this is then to also not respond. Depending on the distribution of the level-k's in play, there may be no response to anticipated event or a muted response. Many of the behavioral mechanisms that we unify have been employed recently in monetary models to explain muted responses to anticipated monetary policy such as Angeletos and Lian (2018), García-Schmidt and Woodford (2019), and Woodford (2019), which incorporate either beauty contest/level-k reasoning, or finite planning horizons arguments, respectively. The unified theory we have laid out and our experimental evidence provides support for these types of approaches.

2 The Model

To develop our unified theory, we consider a univariate nonstochastic model in which agents are differentiated by forecast sophistication. The principal version of the model considers a static framework, which is repeated over time, e.g. the cobweb model as in Muth (1961), with the aggregate variable determined entirely by the expectations of the agents. In this section we first develop the static version of the model, which includes agents with varying levels of forecast sophistication. We then incorporate dynamics via two distinct mechanisms through which agents can improve their forecasts over time. Finally, we present and analyze

the unified model, which joins these two mechanisms.

2.1 The static model

There is a continuum of agents. The aggregate variable at time t, given by y_t, is determined entirely by the expectations of these agents, who are partitioned into a finite number of types. Types are distinguished by sophistication level, which is naturally indexed by the non-negative integers $\mathbb{N}$.[4] For $k \in \mathbb{N}$, the proportion ω_k of agents of type k (i.e. having sophistication k) is referred to as the *weight* associated with agent-type k. The distribution of agents across types is summarized by a *weight system* $\omega = \{\omega_0, \ldots, \omega_M\}$, which is a vector of non-negative real numbers that sums to one, and where M is the number of agent types, which, in our dynamic settings, will typically be endogenously determined and vary over time. We denote by Ω the collection of all possible weight systems as M varies over $\mathbb{N}$. This set, together with its natural topology, will be relevant for some of the analytic work in Section 3.

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

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

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

In our set-up, greater sophistication solely reflects higher order beliefs, as in the level-k framework of Nagel (1995).[5] More concretely we proceed as follows. Agents with *level*-0 *beliefs* hold a common prior and form their forecasts accordingly as $E^0_{t-1}y_t = a$. Agents with higher-order beliefs are assumed to have full knowledge of the model. We recursively define *level-k beliefs* as the beliefs that would be optimal if all other agents used level $k-1$:

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

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

[5] In a macro setting level-k beliefs were used in Evans and Ramey (1992) based on calculation-cost considerations. Recent macroeconomic applications include García-Schmidt and Woodford (2019), who note that their "reflective equilibria" can be interpreted as arising from level-k beliefs with Poisson distribution weights. It also possible to think about more complicated higher order beliefs, where this type of heterogeneity is taken into account by the agents at every level such as in the cognitive hierarchy model in Camerer, Ho and Chong (2004). We view the level of sophistication as an empirical question. Our experimental results show that level-k higher order beliefs well-captures agents' forecasts.

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

The most natural level-0 belief will depend on the model. For example, the level-0 belief may reflect a salient value, as in the guessing game model in Nagel (1995) where this is taken as the midpoint of the range of possible guesses; or, in the cobweb model, the level-0 belief might be determined by the previous equilibrium in a market-setting, before a structural change has occurred, or it may be determined adaptively by looking at past data.

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

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

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

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

2.2 Adaptive dynamics

We define *adaptive dynamics* as corresponding to adaptive learning with fixed level-k weights.[7] Specifically, a weight system ω is taken as fixed and level-0 forecasts $E^0_{t-1} y_t \equiv a_{t-1}$ are assumed to evolve over time in response to observed outcomes. The system under adaptive dynamics is given by

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

where $0 < \phi < 1$. The simple form of the updating rule for level-0 beliefs reflects that our model is univariate and non-stochastic. The parameter ϕ, often called the gain parameter, specifies how much the forecast adjusts in response to the most recent forecast error. The

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

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

time t forecasts a_t can be equivalently written as a geometric average of previous observations with weights $(1-\phi)^i$ on y_{t-i}, for $i \geq 1$.[8]

We could also consider alternative backward-looking rules such as anchor and adjustment rules or trend following rules described by Hommes (2013) and frequently found to well-describe behavior of laboratory participants in LtFEs. These extensions could be introduced into the framework of this paper, but we omit them in our presentation in order to emphasize the novel features of our framework.

2.3 Replicator dynamics

We next consider the possibility that agents revise their depth of reasoning over time based on their past forecast performance. Nagel (1995) and Duffy and Nagel (1997) each explore whether laboratory participants update their depth of reasoning over time in repeated guess-the-average experiments. They find that in general they do not update their reasoning in games with few repetitions - four or fewer - but do appear to do some updating in games of 10 repetitions or more. To capture this sort of updating behavior, we consider the possibility that agents are relatively inattentive to revising their depth of reasoning. Agents who use depth of reasoning that provide a reasonable "forecast" do not revise their reasoning, while only a small proportion of those whose depth of reasoning produces a poor forecasts revise their depth of reasoning in each period with the proportion increasing as forecast performance decreases.

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

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

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

where the left-most "min" is used to break ties.[9]

[8] If y_t, in equation (1), also depended on a white-noise random shock then ϕ would typically be replaced by a time-varying decreasing term such as $\phi_t = 1/t$. In cases where y_t also depends on observable exogenous stochastic shocks, adaptive learning is formulated in terms of recursive least-squares updating.

[9] Note that the existence of the arg min is guaranteed by the fact that if $|\beta| < 1$ then $E^k_{t-1} y_t \to 0$ as $k \to \infty$, and if $|\beta| > 1$ then $E^k_{t-1} y_t \to \infty$ as $k \to \infty$.

Next, assume the presence of a rate function $r : [0, \infty) \to [\delta, 1)$ with $\delta \geq 0$ satisfying $r' > 0$.[10] Finally, let ω_{kt} be the weight of level-k beliefs in period t. The system under replicator dynamics is given by period t according to

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

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

Under the replicator dynamic, when agents revise their depth of reasoning to the optimal k, they do not revise all the way to the eductive solution. Often in models where agents have a choice to be "more" rational, such as in Evans and Ramey (1992), Evans and Ramey (1998) and Brock and Hommes (1997), a cost is imposed to stop all agents from choosing it. In our case, level-k beliefs make it optimal for the agent to not jump to highest level of rationality on offer.

2.4 Unified dynamics

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

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

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

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

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

3 Properties of the unified model

In this section we develop the analytic properties of the unified model. We begin by establishing the available analytic results, and the turn to simulations for additional insights. These insights are aided by some partial analytic results on the dependence of $\hat{k}$ on the feedback parameter β. In the dynamic setting, $\hat{k}$ determines how the depth of reasoning of agents changes over time.

3.1 Stability results

We begin with the stability of the unified model. While the adaptive learning and replicator dynamics can be viewed as special cases of this model, additional insights are available if they are treated separately, and so our first result excludes the cases $\delta_r = 0$ and $\phi = 0$.

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

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

We remark that if $\beta < -1$ then odd levels of reasoning introduce negative feedback while even levels result in positive feedback. These countervailing tendencies can result in interesting and complex outcomes; but they also make $\beta < -1$ difficult to analyze. Some partial results are available under adaptive learning, as discussed below.

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

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

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

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

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

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

The earlier-defined set of all weight systems, Ω, is the disjoint union of these simplexes: $\Omega = \dot{\cup}_n \Delta^n$, where the dot over the union symbol emphasizes that, as subsets of Ω, the Δ^ns are pairwise disjoint. The set Ω inherits a natural topology, sometimes called the *final topology*, from the relative topologies on the Δ^ns: $W \subset \Omega$ is open if and only if $W = \dot{\cup}_n W_n$, with $W_n \subset \Delta^n$ open in Δ^n.[11]

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

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

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

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

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

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

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

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

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

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

3.2 Some results on $\hat{k}$

The behavior of the replicator dynamic is determined by the optimal level of reasoning, $\hat{k}$. To gain intuition for the mechanics of the replicator, in this section we study the dependence of $\hat{k}$ on β for the special case uniform weights. we consider the optimal level of reasoning for a special case of the model.

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

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

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

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

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

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

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

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

3.3 Simulated dynamics of the unified model

To illustrate how convergence is achieved under different specifications of the unified dynamics, we consider a variety of special cases operating under a range of feedback parameters β. In this section, without loss of generality, we set γ at zero, so that $\bar{y} = 0$ (equivalently, the dynamics for y and a may be viewed as in deviation form). We take the parametric form of the rate function for the replicator dynamics to be given by $r(x) = {}^2/_\pi \tan^{-1}(\alpha x)$, with $\alpha > 0$. Finally, all simulations are initialized with $a_0 = 1$ and $\omega_{k0} = {}^1/_4$ for $k = 0, 1, 2, 3$.

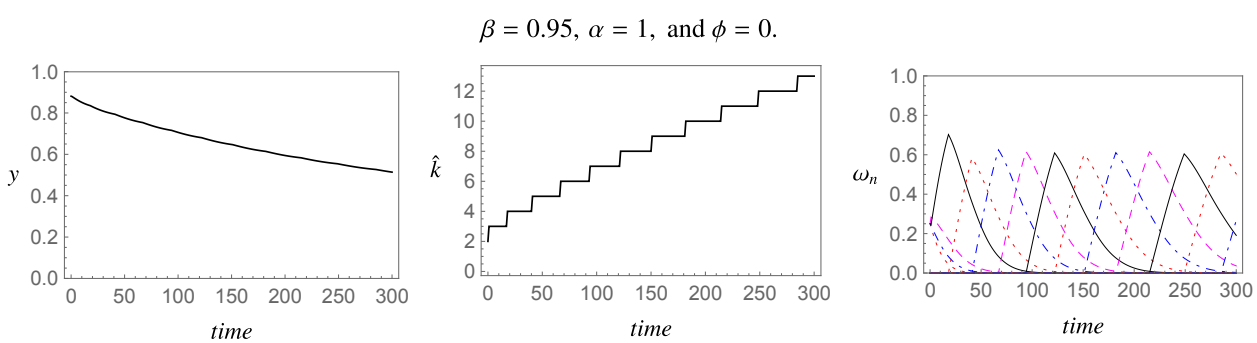

Figure 1: Replicator dynamics with positive feedback.

We start with with the stable positive feedback case $0 < \beta < 1$. Because the results associated with adaptive learning are simply monotonic convergence of a and y to $\bar{y}$, we omit the figure, and thus we begin with replicator dynamics. Figure 1 provides the results with $\phi = 0$ and level-0 expectations are fixed at a_0. Here $\beta = 0.95$ and $\alpha = 1$. The aggregate variable y exhibits monotone convergence to $\bar{y}$, which results from a shift over time in the weight distribution off the k-level forecasts. Note that $\hat{k}$ increases monotonically over time,

and indeed by Theorem 2, $\hat{k} \to \infty$; however, convergence of y to zero is slow: at the end of 300 periods, y has fallen less than half of the distance from the initial value to $\bar{y}$. This slow convergence stands in contrast to the adaptive learning case, i.e. $\delta_r = 0$. For example, when $\phi = 0.1$, by the 300 periods y has fallen by over 90%. This is a reflection of the strong positive feedback value $\beta = 0.95$ and the lack of adaptive learning dynamics.

The right panel of Figure 1 provides the dynamics of agents' weights. As indicated above, $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *dotted, dash-dot, dashed* and *solid*, respectively. As the replicator adds higher forecast levels, the associated paths are graphically identified in an analogous fashion by repeating the styles mod four. We follow this convention throughout this section. As predicted by our theoretical results, we observe lower-level forecasts to gradually fall out of favor and be replaced by higher-level forecasts.

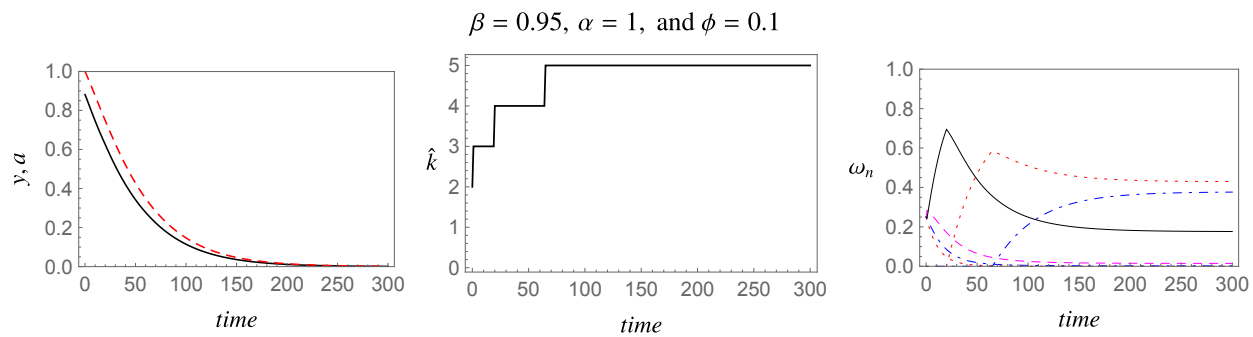

Figure 2: Unified dynamics with positive feedback

Figure 2 gives the results when both adaptive learning and replicator dynamics are operational. Here $\beta = 0.95$, $\phi = 0.1$ and $\alpha = 1$. The left panel now includes the time paths of both y (solid) and a (dashed). Convergence is now much faster, and also faster than the purely adaptive learning case. This reflects the combined role of adaptive learning in shifting the level-0 forecast over time toward $\bar{y}$ and shifting weight toward higher values of k. Note that the weight associated to level k forecasts is increasing if and only if $k = \hat{k}$. For example, we see that around $t = 20$, weight begins to shift to level-4 forecasts, and to level-5 forecasts around $t = 65$. The optimal k appears to stall out at $\hat{k} = 5$ because, as the estimate $a_t \to 0$, higher-level forecasts provide limited to no additional value. Indeed, we have found in other simulations that if the gain is reduced, so that $a_t \to 0$ more slowly, $\hat{k}$ rises further before appearing to plateau.

We now turn to the negative feedback case, with $-1 < \beta < 0$. Again, the results associated with adaptive learning are unexceptional, and indeed, for small ϕ, result in monotonic convergence of a and y. Turning to the replicator dynamics, Figure 3 provides the results for $\beta = -0.5$. First note the non-monotonic behavior of y: the left panel, scaled differently in this figure, shows oscillatory convergence of y induced by the negative feedback. The behavior of $\hat{k}$ reflects these oscillations: when y crosses zero, $\hat{k}$ rises sharply to drive down (in magnitude) the optimal forecast $\beta^{\hat{k}}$. There are, in fact, three times that y crosses zero

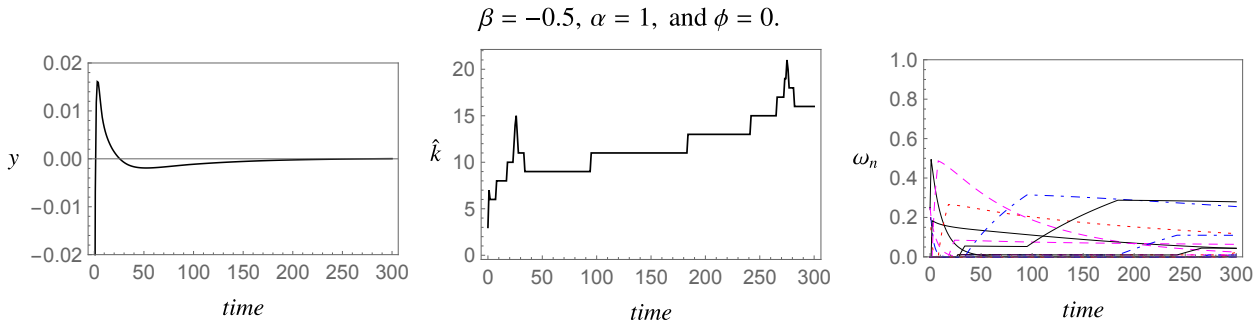

Figure 3: Replicator dynamics with negative feedback.

during the observation period, at $t = 2, 26$, and 276, which are associated with spikes in $\hat{k}$, particularly visible in the latter two cases. Finally, we note that by Theorem 2, $\hat{k} \to \infty$, however, unlike the positive feedback case, here this convergence is not monotone.

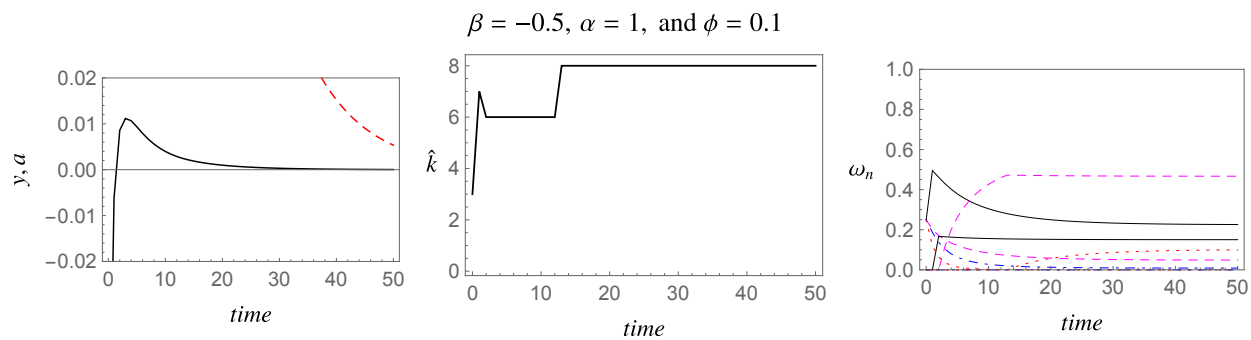

Figure 4: Unified dynamics with negative feedback.

Figure 4 gives the results when both adaptive and replicator dynamics are operational. Because adaptive learning drives level-0 forecasts to zero there is faster convergence, with weaker oscillatory behavior, than in the replicator-only case.

Finally, we turn to the case in which $\beta < -1$. We remark that, in this case, $\bar{y}$ is not stable under eductive learning as shown in Guesnerie (1992): if all agents are fully rational and have common knowledge of the structure they are unable to coordinate on the REE. However, as indicated by Corollary 1, when $\beta < -1$ the REE is stable under adaptive dynamics provided the gain is sufficiently small.

In the replicator-only case, the dynamics can be unstable or can exhibit complex behavior. For example, Figure 5 provides a simulation with $\beta = -2.0$ and $\alpha = 0.05$. Note that $\hat{k}$ oscillates between 0 and 1, which drives ω_{nt} to zero for $n \geq 2$. For viewing convenience, the right panel, which plots the time paths of these weights, does not employ the style convention used above; instead, the higher oscillating path corresponds to ω_0 and the lower one to ω_1. The evolution of y appears to converge to an 11-cycle, which, we observe, is not centered

at zero.[12] Figure 6 exhibits the corresponding simulation with unified dynamics in which $\phi = 0.1.$, and uses the same style convention as Figure 5. The addition of adaptive learning dynamics pushes level-0 expectations towards zero, which when combined with replicator dynamics leads to rapid convergence to the REE.[13]

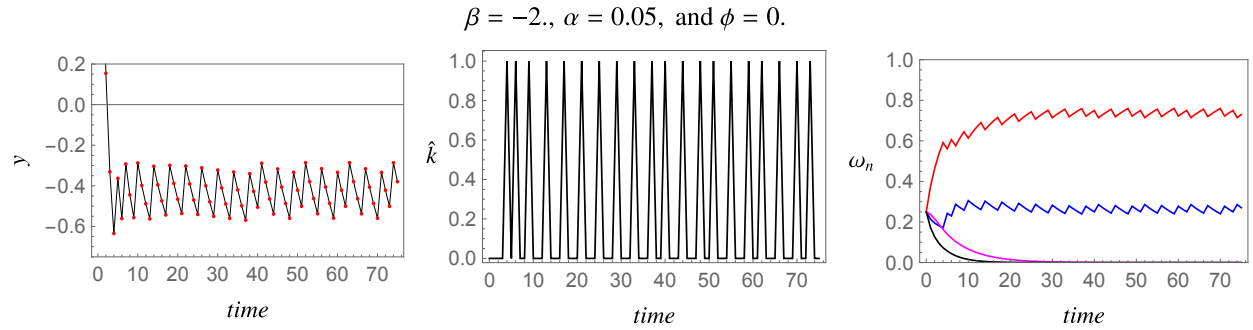

Figure 5: Replicator dynamics with large negative feedback.

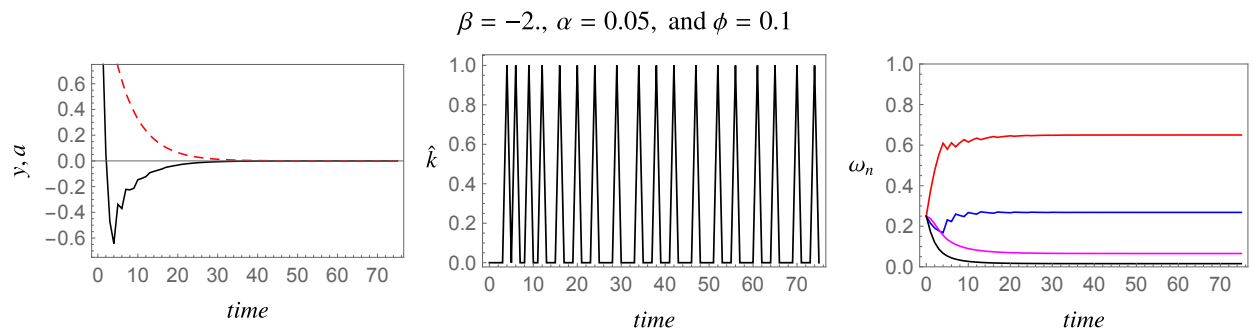

Figure 6: Unified dynamics with large negative feedback.

4 Learning-to-Forecast Experiment

Our model has a number of testable implications that can be profitably examined in a laboratory setting. We do this using a standard LtFE experiment following Hommes et al. (2007), Hommes (2011), and Bao et al. (2012), in which groups of participants are asked to forecast prices in a simple market environment, and are paid based on the accuracy of their forecasts.

The experimental market we consider takes same form as the one used to describe the unified model:

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

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

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

where p_t is the market price, $\hat{E}_{t-1}p_t$ is the average price forecast across participants, $\beta < 1$, and we introduce a small white noise shock ϵ_t as is standard in these experimental settings.[14]

The testable implications of our model involve how stability, speed of convergence, and depth of reasoning depend on the expectational feedback parameter β. One contribution of our experiments is to examine participants' responses to announced structural change. Specific, testable hypotheses are given and discussed below. To address the hypotheses, we adopt a 3×3 design where the treatment variables are (1) the strength of the feedback of expectations ($T\#$) and (2) the timing and size of the announcement of a structural changes in γ ($A\#$).

The feedback treatments and their predicted stability properties are

T1: $\beta = -0.9$. Theorems 1 - 3 imply asymptotic stability of the REE for any specification of the unified dynamic, i.e. for $\phi \in [0, 1]$ and $\delta_r \in \{0, 1\}$. The simulations of Section 3.3 suggest rapid and possibly oscillatory convergence.

T2: $\beta = -2$. Formal results obtain from Theorem 3, when $\delta_r = 0$. For a given weight system, the Theorem implies asymptotic stability for sufficiently small gain ϕ. Also, for given ϕ, the Theorem implies a large set of weight systems resulting in asymptotic stability and a large set yielding instability. Together, these results suggest that asymptotic coordination on the REE is challenging under unified dynamics. The simulations indicate that under the replicator, stable cycles can emerge, while under the unified dynamic, non-monotonic convergence may obtain.

T3: $\beta = 0.5$. Theorems 1 - 3 imply asymptotic stability of the REE for any specification of the unified dynamic. The simulations of Section 3.3 suggest monotonic converge at a rate that is slower the negative feedback counterpart.

The announcements treatments are

A1: γ increases from 60 to 90 in period 50

A2: γ increases from 60 to 90 in period 20 and remains at the new value for the remainder of the experiment

A3: γ increases from 60 to 90 in period 20 and decreases to 45 in period 45.

4.1 Simulations of the experiments

We now turn to simulations of the model to illustrate the predicted individual forecasting behavior of agents and the dynamics of the endogenous variable p_t for different values of β

[14] We use the same random process for the white noise shock as used by Bao and Duffy (2016).

Figure 7: Positive feedback case $\beta = 0.5$

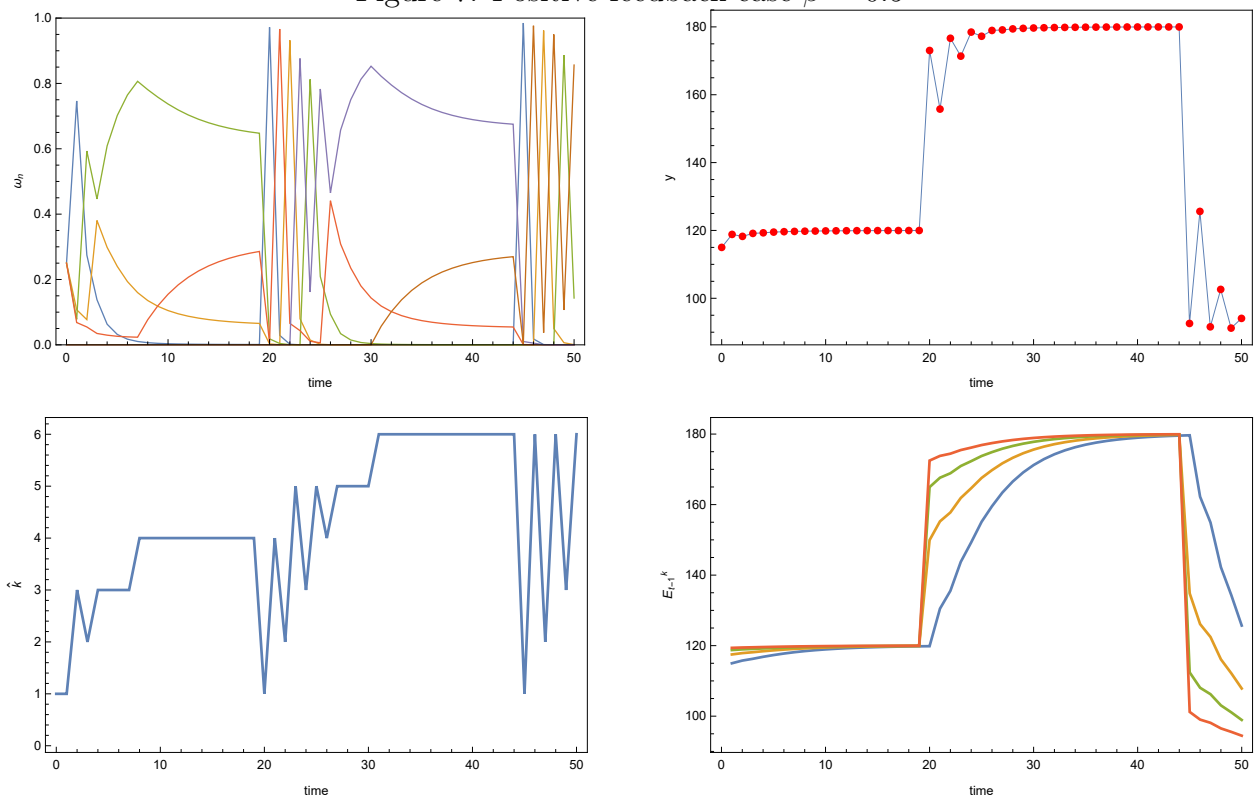

Notes: The four panels of the figure moving clockwise show the evolution of the mass of agents that choose a given predictor, the evolution of y_t over time, the prediction of the levels 0, 1, 2, and 3 forecasts given the current state of the economy, and the optimal $\hat{k}$.

and in response to announced structural change to γ. To make these predictions comparable to our experiment, we impose a non-negativity constraint for the endogenous variable p_t. The floor is assumed known by the agents and incorporated into the level-k beliefs.

Our stability propositions identify two separate regimes of interest $|\beta| < 1$ and $\beta < -1$ as well as two subcases within each. We explore the same simulation for each relevant case. Specifically, mirroring our experimental setup with the A3 announcement treatment, we consider a 50-period simulation with $\gamma = 60$ for $t < 20$, $\gamma = 90$ for $20 \leq t < 45$, and $\gamma = 45$ for $t \geq 45$. We initialize the adaptive beliefs, level-0, to be below the PFE in time period $t = 0$ to begin each simulation.

Figure 7 shows the simulation for a positive feedback case, where $\beta = 0.5$, $\phi = 0.2$, $\alpha = 0.5$, and assuming that agents are evenly distributed across levels-0, 1, 2, and 3 in the initial period. The four panels of the figure moving clockwise show the evolution of the mass of agents that choose a given predictor, the evolution of p_t over time, the prediction of the level-0, 1, 2, and 3 forecasts given the current state of the economy, and the optimal $\hat{k}$. The value of p_t converges quite quickly to a value close to the PFE. During this time the

replicator redistributes the mass of forecasters to a succession of different optimal k's. First, level-1 is the best forecasts, then level-3 and then 4. This change correspond to rise and fall of the weights in the top left panel above.

When the first announcement occurs, significant heterogeneity is seen across the different k types. Forecasters with higher levels of k expect higher values of p_t while level-0 forecasts ignore the information entirely (blue line). Because of this heterogeneity, the actual price does not move as high as the high level-k agents expected. Lower levels of k result in better accuracy, which drives down the optimal $\hat{k}$. This introduces some volatility before convergence resumes and $\hat{k}$ continues higher. The same dynamic occurs again for the second announcement.

The optimal level of k advances slowly in this case because of the positive feedback. The endogenous variable moves in the direction of the forecasts but less than one-for-one. Therefore, if agents jump to too large of a k, then their forecast can outrun the actual realization of p_t. This is why it is optimal to choose low levels of k for long periods of time and why convergence to common knowledge of rationality only occurs in the limit.

Figure 8: Positive feedback case $\beta = -0.9$

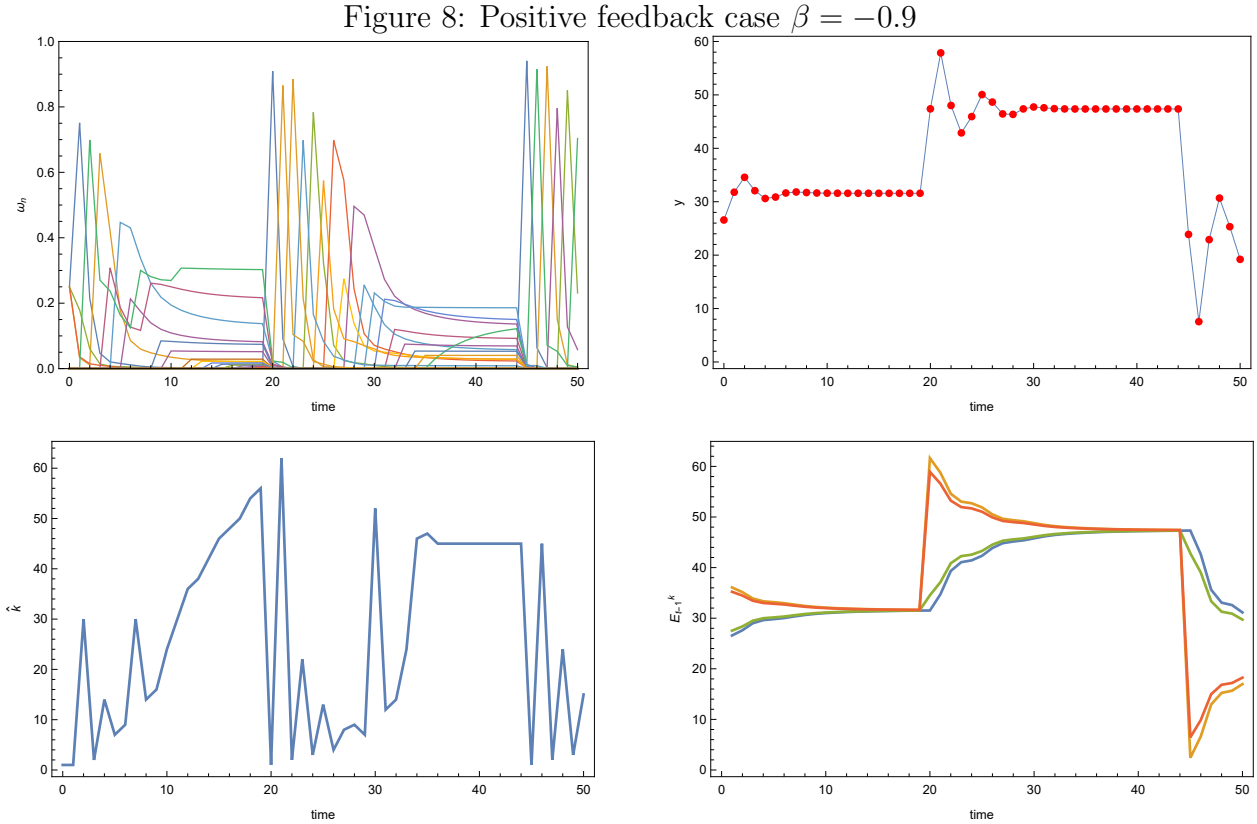

Notes: The four panels of the figure moving clockwise show the evolution of the mass of agents that choose a given predictor, the evolution of y_t over time, the prediction of the levels 0, 1, 2, and 3 forecasts given the current state of the economy, and the optimal $\hat{k}$.

Figure 8 shows the simulation for $\beta = -0.9$ with all the remaining parameters the same

as before. Again, as predicted by our propositions, there is convergence toward the PFE with large $\hat{k}$ values observed. Negative feedback provides clearer signals to the agents about what level of k is best since p_t often moves in the opposite direction of an agents' forecasts. This leads to a faster rise in the optimal k in this case.

Figure 9: Positive feedback case $\beta = -2$

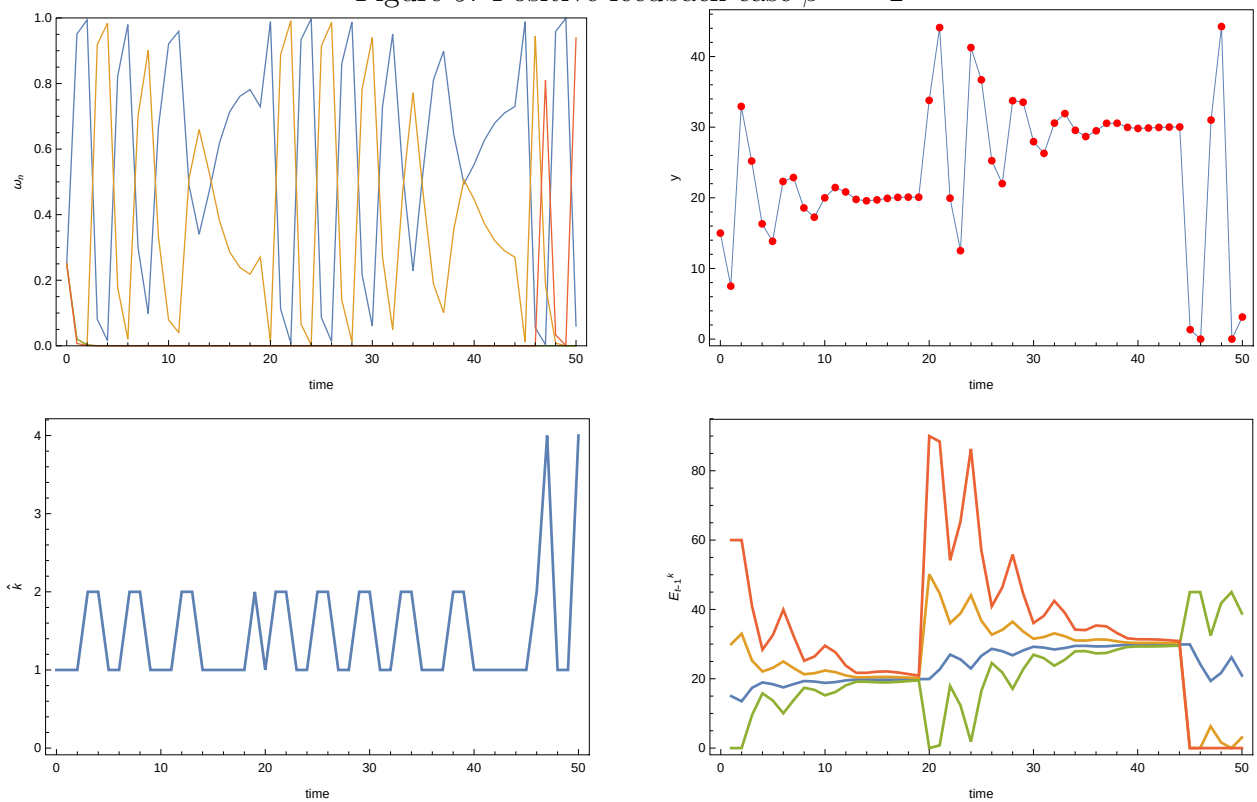

Notes: The four panels of the figure moving clockwise show the evolution of the mass of agents that choose a given predictor, the evolution of y_t over time, the prediction of the levels 0, 1, 2, and 3 forecasts given the current state of the economy, and the optimal $\hat{k}$.

Figure 9 shows the simulation for $\beta = -2$ with all the remaining parameters the same as before. This case is not covered by our stability propositions. For this particular parameterization, though, we still find that convergence to the PFE occurs relatively quickly. In fact, the dynamics of p_t, the individual level-k forecasts, and their responses to the announcements all behave similarly to what was observed in the previous simulation.

The primary difference observed in this case is in the evolution of $\hat{k}$. The optimal k remains low for most of the simulation. This is because the high negative feedback causes predictions to diverge toward the non-negativity constraint. Once the constraint is hit, successive level-k forecasts ping-pong from this constraint to the maximum price implied by a zero price in the previous period.

It also important to note again that convergence is not guaranteed when $\beta < -1$. As

shown in Section 3, convergence depends on relative size of ϕ and α and in a nonlinear way. For example, it is possible to destabilize the simulation shown by both increasing or decreasing ϕ from the value chosen in Figure 9.

4.2 Testable hypotheses of the unified theory

The 3×3 design allows us to test the following hypotheses based the simulated properties of the unified model, our propositions, and on the results of Bao and Duffy (2016).

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

Bao and Duffy do not find evidence of pure eductive reasoning. But they do find different behavior in cases when expectational feedback satisfies eductive stability relative to when it does not.[15] This corresponds to treatments T1 and T3 versus T2 in our experiment. Specifically, Bao and Duffy find that convergence to the REE happens more slowly or not at all when Eductive stability is not satisfied. Our unified theory makes the same predictions. This is seen by comparing, for example, Figures 3 to **??** or Figures 8 to 9.

Hypothesis 2: Participant's predictions in announcement periods in treatments A1 - A3 are well described by level - k deductions for all treatments T1 - T3.

The key treatment in this experiment is the introduction of credible announcements about structural change in a full information environment. Our A1 - A3 treatments are akin to inserting a beauty contest games into an otherwise standard LtFE. The advantage of this setup is that we can precisely identify whether agents form high order beliefs following level-k deductions by observing their predictions when announcements are made. Identification is obtained because we observe the history of prices that the participants have observed allowing us to infer the value of a level-0 forecast in accordance to the unified theory. All relevant level-k forecasts can then be defined from this reference point and compared to the laboratory participants' predictions.

Hypothesis 3: Participants revise their level of reasoning in response to losses and toward the most recent optimal predictor.

A defining feature of the unified theory is the replicator dynamic that governs how agents revise their depth of reasoning in response to poor forecast performance. The unified theory posits that not all agents revise in every period but those that do, revise towards the optimal predictor based on last periods price. This state dependence makes it so that revisions

[15] When the number of participants in a market is finite, the eductive stability condition is relaxed to $-N/(N-1) < \beta < 1$. Therefore, the appropriate condition for our experiment is $-6/5 < \beta < 1$.

to the depth of reasoning may be non-monotonic in some instances. Therefore, following announcements, we expect revisions downward for those agents who experience large forecast errors as shown in Figures 7, 8, and 9.

Hypothesis 4: The average depth of reasoning is increasing for the T1 and the T3 treatments over time but not for the T2 treatments.

Proposition 4 provides an extension to the Eductive stability conditions of Guesnerie (1992). We prove that convergence under the replicator dynamic only is obtained when $|\beta| < 1$. When $\beta < -1$ non-convergence, or cycling behavior are all possible depending on the parameters of the model. Therefore, we expect to have different behavior for revisions of the depth of reasoning for T1 and T3 compared to T2. Specifically, the average level of reasoning in T1 and T3 treatments should increase over time whereas it should not necessarily in T2.

4.3 Experiment description

The experiment is conducted in a computer based market programmed in oTree.[16] Laboratory participants are told that they are acting as expert advisers to firms that produce widgets. Participants are then led through a tutorial that describes the market environment including the demand and supply equations that govern price, the fact that price depends on the average expected price of all advisers in the market, and that prices are subject to small white noise shocks. The market setup is as close a recreation as possible to that of Bao and Duffy (2016) that accommodates our novel elements.

The market stories differ for the positive (T3) and negative feedback cases (T1 and T2). The latter describes the normal cobweb setup of perfect competition among firms that face convex costs of production of a non-storable good. The former introduces the notion that a widget is a Veblen good with upward sloping demand. In each case, the fact that there is negative or positive feedback is stated multiple times for the participants and even printed in bold on a paper versions of the instructions that participants receive with the following wording:

> *"KEY POINT: The market has positive feedback. Therefore, if the average price forecast is high, the market price will be high. And, if the average price forecast is low, then the market price will be low."*

The negative case is stated similarly. The tutorial and printed instructions are available in Appendix D.

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

Figure 10: Screenshots of experimental market

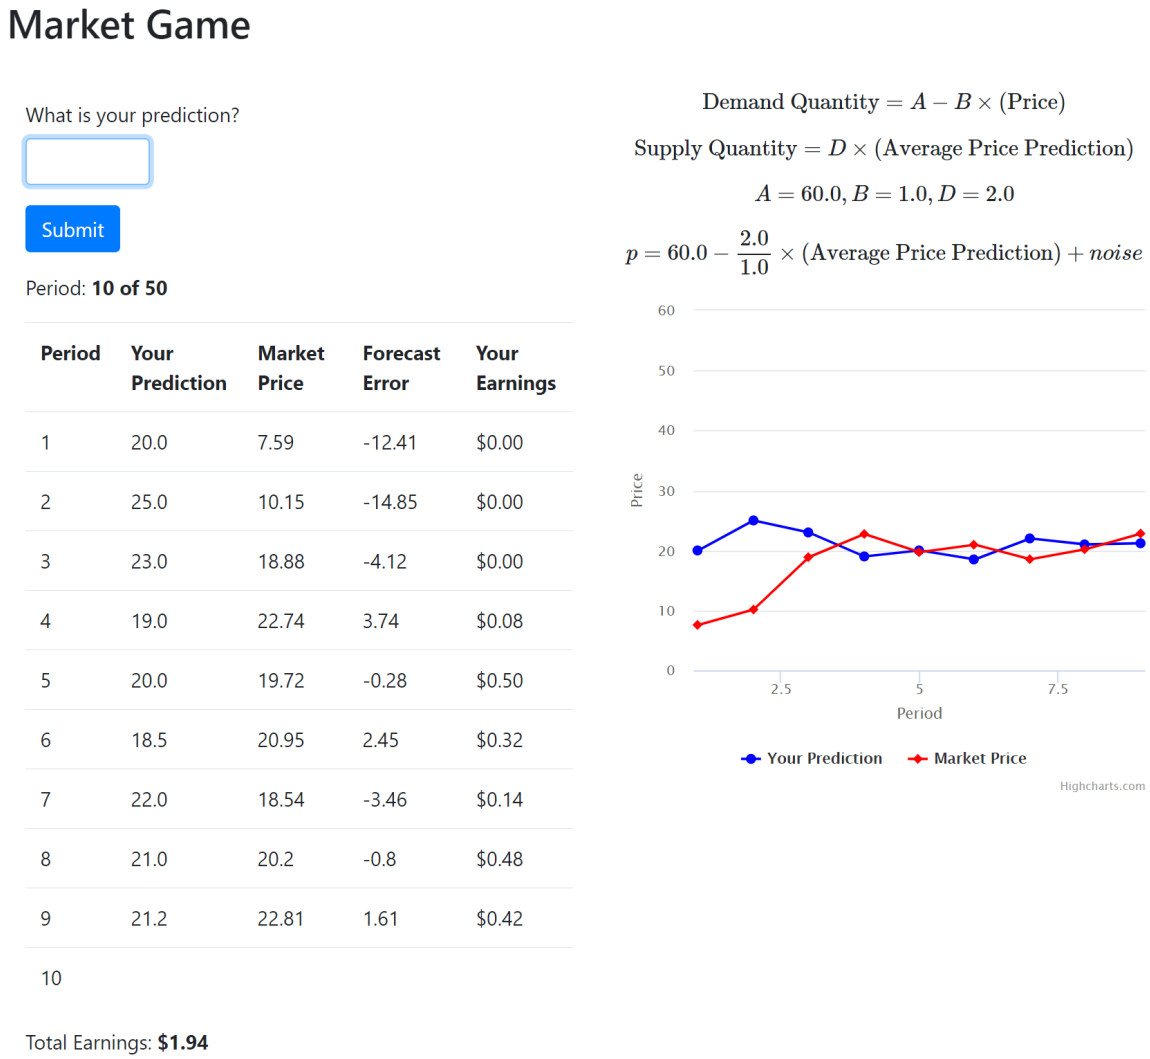

Period	Your Prediction	Market Price	Forecast Error	Your Earnings
1	20.0	7.59	-12.41	$0.00
2	25.0	10.15	-14.85	$0.00
3	23.0	18.88	-4.12	$0.00
4	19.0	22.74	3.74	$0.08
5	20.0	19.72	-0.28	$0.50
6	18.5	20.95	2.45	$0.32
7	22.0	18.54	-3.46	$0.14
8	21.0	20.2	-0.8	$0.48
9	21.2	22.81	1.61	$0.42
10				

Figure (10) shows the graphic user interface (GUI) for the T2 treatment that participants interacted with during the experiment.[17] The market information is shown in the top right of the screenshot. We checked for comprehension of the market environment with a version of the following question at the end of the tutorial:

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

[17] A positive feedback example is given in the appendix

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

we get that price (p) is

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

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

A worked version of this problem with different numbers was also provided on the printed instruction sheet. Participants were not able to continue with the experiment until the question was answered correctly.[18] The question was designed to make sure that participants knew how to use the equations without teaching them to solve for the REE. We believe that this question also provides us with a reasonable degree of certainty that participants understand that the outcomes depend on the average price forecasts. The GUI also provided participants with a time series plot of the price and their predictions (bottom right) and past prices, predictions, forecast errors, and period earnings in table form (left-hand side).

The period-by-period payoff function for the participants is given by

$$\text{payment}_t = 0.50 - 0.03\,(p_t - p_t^e)^2$$

where p is the actual market price, p_t^e is their prediction, and 0.50 and 0.03 are measured in cents. Negative quantities receive zero cents. The function is presented and explained to participants as part of the tutorial and is the same across treatments.

This payoff function is a slight departure from previous LtFEs in that it is directly measured in hard currency rather than points and it is stricter than many comparable studies. Forecasts must be within 4 units of the actual price or the participant earns nothing. We chose this specification to give participants a relatively high incentive to be precise in their predictions when confronted with announcements and to give more salience to losses. Previous studies have employed point systems that compensate more generously for poor forecasts. For example, in Bao and Duffy (2016) participants needed to be within 7 units to earn points, which ranged from zero to 1300.

Announcements were made to participants using a pop-up box at the beginning of the period in which they occur. The popup box would describe the change in parameters. Participants were required to close the pop-up before they could proceed. The announcement

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

would also appear, highlighted in red, across the top of the screen. The information in the top right corner would also reflect the change. Each participant played 50 rounds. Afterwards, participants were surveyed on the strategy they employed, what strategy they believed others employed, and what information they found most useful.

5 Experimental results

This section summarized the experimental results. We provide an overview of the results and sample in the first subsection and then provide evidence relating to each individual hypothesis.

5.1 Overview of the data

Table 1 reports summary statistics for participation and pay in the experiment. In total, 372 individuals participated in 62 experimental markets. All T1 and T2 treatments were conducted in May and June of 2018 at the UNSW Sydney BizLab. Two sessions for each treatment were scheduled with the aim of testing at most five markets in each session. Participant no-shows account for the different number of markets across the treatments.[19] All T3 treatments were conducted in March of 2019 at the University of Sydney's Experimental Lab. Eight sessions were held with the aim of testing at most four markets in each session.[20] Again, no-shows account for different number of experimental markets across treatments.

In addition to performance pay, subjects received a $5 show-up fee in the T1 and T3 treatments and a $10 show-up fee in the T2 treatment. The higher show-up fee in the T2 treatments was given to compensate participants for lower earning that we expected would result from difficulties in coordination. Sessions lasted on average 80 minutes including verbal instructions. Groups that exceeded 100 minutes were paid additional compensation to make sure that average payouts were at least $15 (AUD) an hour to satisfy ethics requirements. Participants were not told in advance that they would receive this additional compensation.[21]

Figure 11 shows the average price observed across all treatments relative to the PFE price. Figure 12 shows the individual price predictions for all individuals with outliers indicated by X's. The individual forecasts illustrate both the diversity and uniformity that can occur depending on the expectational feedback in the market. As predicted by the simulations shown in Section 3.3, all the $|\beta| < 1$ cases show convergence to the REE initially and

[19] Subjects were recruited using ORSEE (Greiner (2015)). A3 treatments were specifically scheduled at times of day when participant turn-up rates are historically high since the total number of individual responses to announcements is a key variable of interest.

[20] The different session sizes at the University of Sydney reflect lab capacity constraints due to equipment issues and subject recruitment limitations.

[21] An announcement was made to participants that they would be compensated for the extra time only when the time threshold was reached.

Table 1: Summary statistics

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

Notes: Participants also received a $5 show-up fee in T1 and T3 treatments and $10 show-up fee in T2 treatments. Pay efficiency is the total pay divided by the maximum pay of $25.

after the announcements, whereas both convergence and non-convergence is observed when $\beta < -1$.

We observed more outliers in individual predictions in this study than were observed, for example, in Bao and Duffy (2016). However, we also have more than double the participants. Some outliers are easily explained as "fat finger" errors where an extra zero is added to a forecast. Others reflect participants with a penchant for anarchy who consistently typed in nonsensical forecasts. For the most part these outliers did not have a significant effect overall on market convergence as we will show.

5.2 Results on convergence

Bao and Duffy's key experimental treatment is a comparison between T1 relative T2. Under the assumption that laboratory participants are purely adaptive learners, all markets should converge to the equilibrium price. However, under the assumption that participants are eductive, coordination should be obtained in treatments T1 but not in T2. Convergence under pure eductive reasoning should also occur from period one on. Likewise, our T3 treatment, which was not considered by Bao and Duffy, has the same eductive stability predictions as T1. Under the assumption of the unified model, convergence should occur under T1 and T3 but may not occur under T2 for some experimental groups.

We have 111 market observations from which to test for convergence spread more-or-less evenly across all feedback treatments. We gain these extra sample markets by looking at convergence both before and after the first announcements in the A2 and A3 treatments. Because of the random noise term, a natural way to judge convergence is to look at markets prices that lie within some tolerance of the steady state price. Following Bao and Duffy, we set this tolerance at $\pm$ 3. We define a market as converged when we first observe two consecutive periods with the price within 3 of the steady state price.

Figure 13 shows the percentage of markets that we classify as converged in each period for the three treatment pairs. Table 2 reports the summary statistics for convergence. We report

Figure 11: Average market price relative to REE

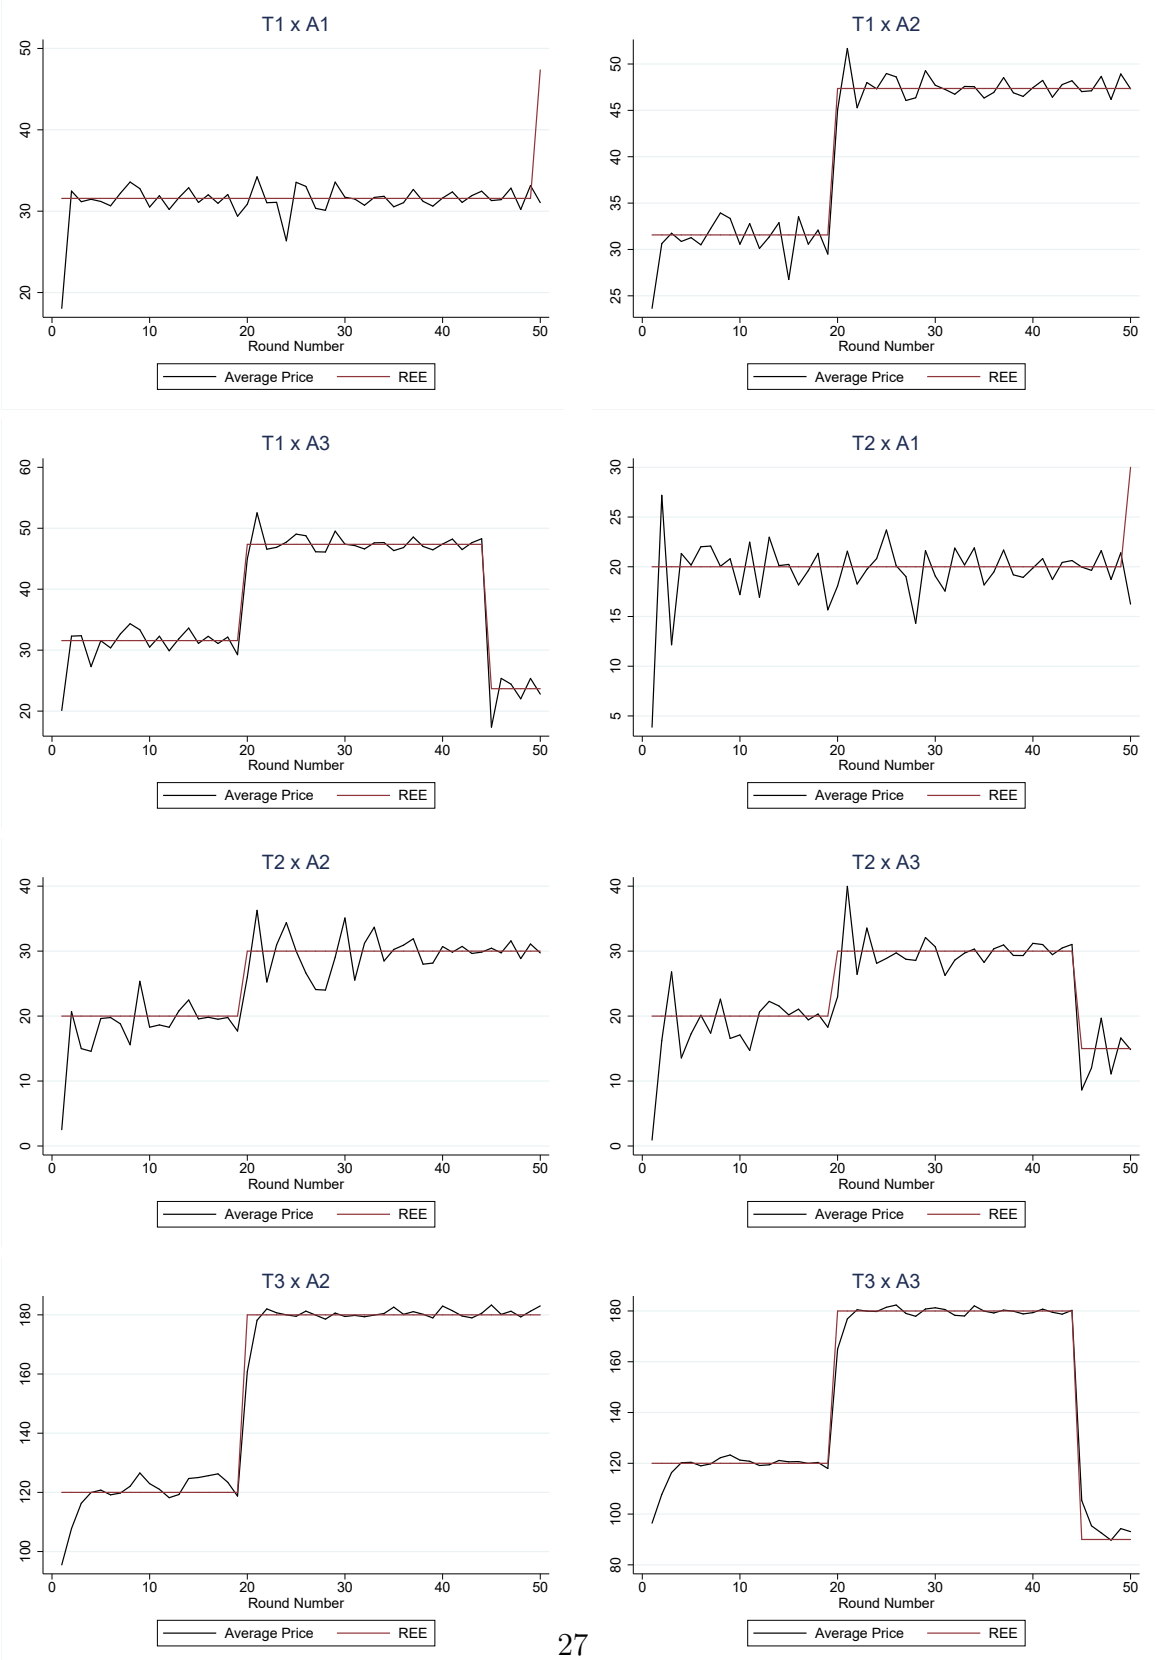

Figure 12: Individual participant predictions

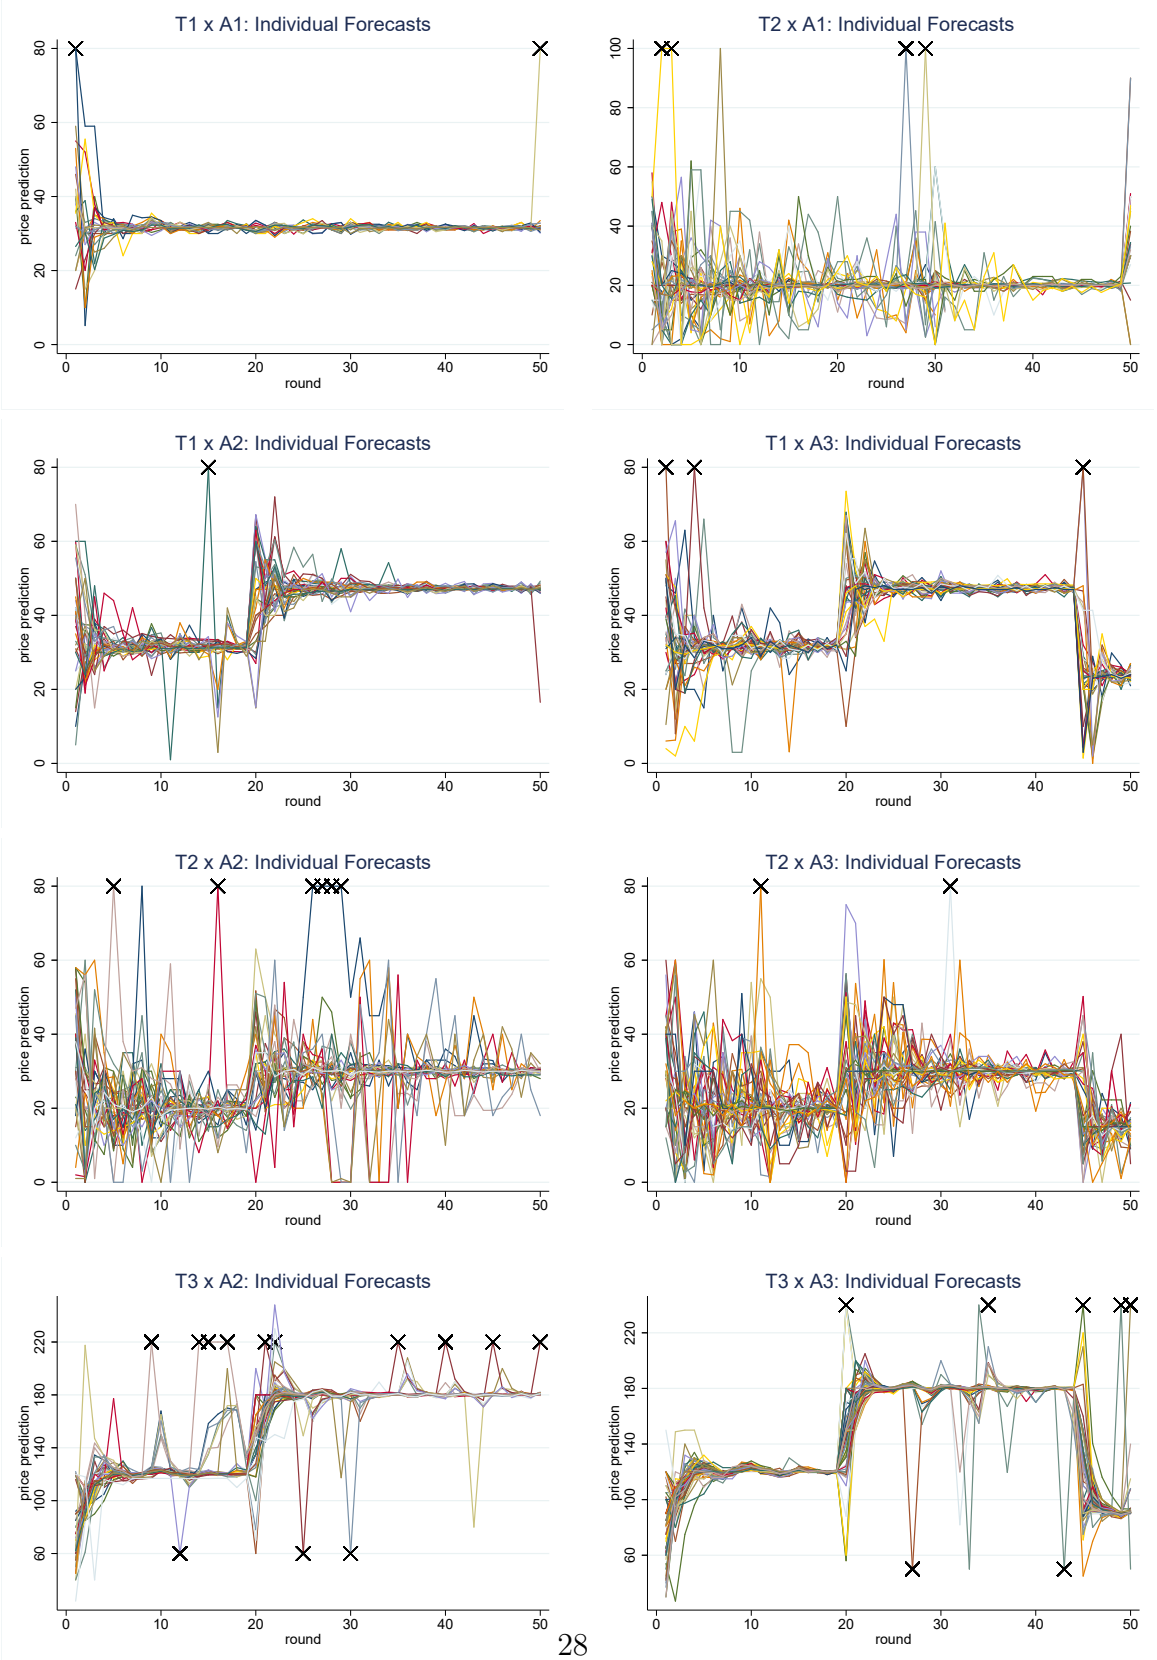

Notes: The 'X's denote forecasts that are larger than the top axis shown in the graph. The maximum value the program would allow a participant to predict is 500.

Figure 13: Convergence over time

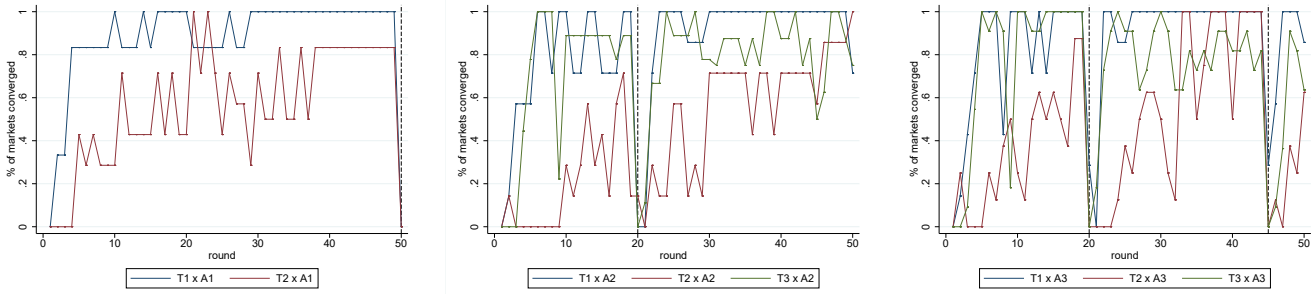

Notes: A market is classified as converged when we observe two consecutive quarters of the price within 3 of the steady state price.

the median period for which we first observe convergence, the max period if converged, the count of markets that never converge, and finally the median percent of remaining periods (PRP) that we observe the market price within the band following the first classification of convergence. The median PRP within $\pm$ 3 shows that the majority of markets remain converged once they are first classified. We only observe a complete failure to converge in the T2 treated markets ($\beta = -2$) with 3 out of 37 markets failing to convergence. This is consistent with the predictions of the unified model and the findings of Bao and Duffy. We also find lower rates of markets staying within the band after initial convergence is recorded in this case.[22]

The median number of periods observed for convergence falls for all treatments after the first announcement. Mann-Whitney tests of whether the time-to-convergence for the initial period is from the same population as the time to convergence for the second period are rejected at standard significance levels for T1 and T3. But we fail to reject the null for the T2 treatments. This is consistent with the unified model, where the errors experienced after an announcement should cause upward revision to an agent's depth of reasoning leading to faster convergence in the T1 and T3 cases but not in the T2 cases.

Another way to assess convergence is to look at average earning. Figure 14 shows the average earning from period 1-19 plotted against period 20 - 38 for treatments with an announcement in period 20. Recall that in order to earn money for a forecast it must be within $\pm$ 4 of the actual price. Therefore, if earnings are high on average, the market must have converged. The maximum earning that a participant can obtain is 50 cents. However, because of the random noise component it is rare for participants to earn exactly that amount. Therefore, when the average earnings are around 40 cents it is a good indicator that a market has converged. We can see that convergence appears to be quicker in all cases after the first announcement then in the beginning periods, which is consistent with increasing depth of

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

Table 2: Convergence results

Treatments	# Markets	Time to converge Median	Time to converge Max	# Markets never converged	Median PRP within ± 3
T1 (per. 1 - 20)	20	4	7	0	100%
(per. 20 - 44)	14	3*,†	4	0	100%
T2 (per. 1 - 20)	22	7	19	2	88%
(per. 20 - 44)	15	5.5	14	1	93%
T3 (per. 1 - 20)	20	5	7	0	93%
(per. 20 - 44)	20	3*,†	9	0	100%

Notes: PRP stands for percent of remaining periods. We include the A1 treatment results with the per. 1-20 results but only declare non-convergence if it is never observed in the full 50 periods in these cases.
∗ Indicates a rejection of the null hypothesis that the sample of market convergence times for the initial period is from the same population as the second at the 5% significance level using a Mann-Whitney test.
† Indicates a rejection of the null hypothesis of the equality of medians at the 5% significance level using a non-parametric k-sample test.

reasoning in the T1 and T3 treatments. The explanation for the faster convergence in the T2 case is partly explained by more participants choosing forecasts that are consistent with adaptive strategies, i.e. revising down their depth of reasoning.

Figure 14: Average period earnings before and after the first announcements

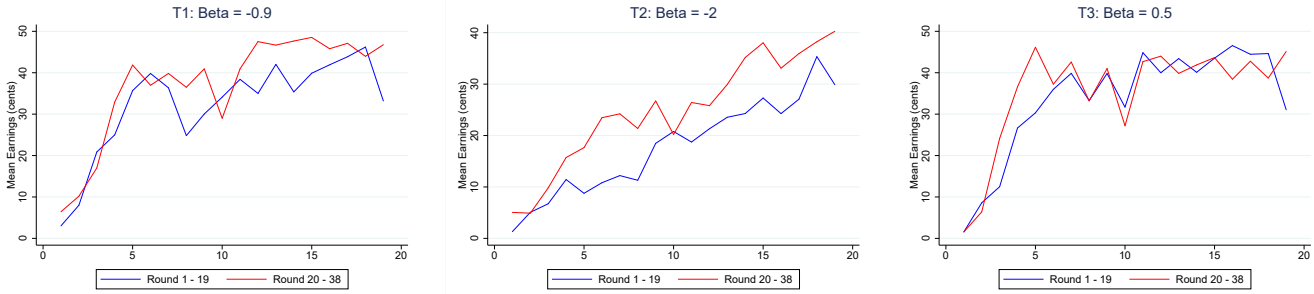

Notes: These graphs show average earnings for periods 1-19 and periods 20-38 for treatments when there is an announcement in period 20. The graphs illustrate faster convergence after the announcement. T-tests of the disaggregated series show statistically different rates of convergence for all earnings treatments.

Comparing across the feedback treatments in Figure 13 we can also see some heterogeneity between T1 and T3. Fehr and Tyran (2008) argue that convergence should be faster in the T1 environment than in T3. The unified model concurs. Comparing the distribution of first period of convergence dates from T1 and T3 for periods 1 - 20, we find no statistical difference using a two-sample Kolmogorov-Smirnov test for equality of distribution functions. But we do find a statistically significance difference following the second announcements (p-value

0.004) with the T3 taking longer to converge as predicted. Pooling all observations likewise results in a statistically different rate of convergence (p-value 0.058) with T3 converging slower than T1 treated markets.

5.3 Evidence for level-k reasoning

The unified model makes strong predictions about how laboratory participants' forecasts should behave following an announcement. Each announcement is preceded by at least 19 periods of play, which is ample time for the market to converge to the steady state price as we have seen. Therefore, the natural definition of a level-0 forecast is the steady state price. From this price, we can define all other level-k forecasts with which we can compare to the subjects' observed forecasts.

Figure 15 shows histograms of the forecasts chosen by agents in the A2 and A3 treatments pooled by the feedback in the market. We have indicated where the level-k forecasts are for each case assuming that the level-0 agents' forecast is the steady state price prior to the announcement. Clearly, the majority of forecasts are centered around the level-k predictions.[23]

To be more concrete, we define a prediction as being consistent with level-k forecast if it is within $\pm$ 3 of theoretical level-k prediction implied by a level-0 forecast at the steady state price prior to the announcement. Likewise, we classify a subject's forecast as level-0 if it is within 3 of the steady state price. Table 3 shows the classification results under this definition. We find that approximately 70% of laboratory participants forecasts can be classified as level-0, 1, 2, 3, or REE combining all announcement treatments excluding the first period.

In addition, we also find qualitative support for level-k reasoning in the exit survey results. In the exit survey, participants were asked about what information was most important in making their decisions. The most common ranking was past prices, their expectation of average price, the equations, their past forecasts, their past forecast errors. They were similarly asked what they believed other people in the markets thought was the most important information when making a decision. In that ranking, the expectation of the average price was overall rated as less important and past forecasts and forecast errors as more important to others players than they were rated for their own decisions. This is consistent with players viewing their opponents as less sophisticated than they view themselves.[24]

[23]The only outlier in Figure 15 is in the T3 x A2 and A3 treatment figure. Here there is a mass of forecasts around $p = 135$. This is the forecast a person would choose if she believed that the other average forecast would be exactly equal to the new intercept value, $\gamma = 90$, i.e. $90 + 1/2^* 90 = 135$.

[24]Summary results from the exit survey are available in appendix E.

Table 3: Classification of subject types

	Announcements Periods			
Type	1	20	45	50
REE	16.7%	23.8%	40.4%	29.6%
Level 0	-	9.9%	6.4%	3.7%
Level k (1, 2, 3)	-	37.1%	23.1%	33.3%
% Classified	16.7%	71.1%	69.9%	66.7%
N	372	294	156	54

Notes: We define an individual prediction to a level-k or REE forecast if it is within ± 3 of the model implied value assuming that level 0 has converged to the steady state price in period prior to the announcement. Since we do not have a prior period for period 1, we only can characterize REE predictions with certainty. We are also missing 24 observations for announcement date 50 because of a coding error.

5.4 Evidence of the replicator

The replicator has three important assumptions. First, it assumes that not every agent will update their forecast in every period. Second, the agents who do are on average those who experience large errors in the most recent period. And finally, the agent's choice of a new strategy is based on a counterfactual exercises where the performances of alternative strategies is assessed by looking at how they would have performed in the previous period.

The announcement in the experiment offer a nice setting in which to test all three of these assumptions. The differences in level-k forecasts are largest given an announced change. This makes identifying the strategy chosen by each participant much easier compared to when a market is converged and most strategies produce similar forecasts. Likewise, the announcement causes almost all participants to experience an increase in their forecast error and a decrease in earnings (see Figure 14). Therefore, by looking at changes in the announcement period and the period following the announcement it is straightforward to assess who has changed strategy, how the strategy choice compares to the strategy in the announcement period, and whether those who change experienced larger errors.

For this exercise, we consider an alternative way to classify which strategy a participant is using that provides a universal identification of subject's forecasts as either level-0, 1, 2, 3, or REE. We do this by constructing participant specific counterfactual forecasts for level-0, 1, 2, 3, and REE that we can compare to the actual prediction made by each participant.

The counterfactual forecast that is closest to the actual forecast as measured by squared error in a given period classifies the participant's forecast type. We construct counterfactual forecasts using Equation (**??**) with $\phi = 0.5$ for each individual market using the observed market price to define an adaptive learning level-0 forecast.[25] Then, based on this level-0 forecast, we construct the implied level-1, 2, and 3 forecasts in each period. We find that approximately 85% of all subject forecasts and 75% of subject forecasts during announcement periods have a root mean squared error of three or less compared to this counterfactual.[26] The higher rate of identification we achieve here is due to the fact that the level-k forecasts are defined off of market specific level-0 forecasts rather than steady state price as in the previous section.

In addition to providing a classification to all participants, classifying participant forecasts in this way provides a benchmark to test the third assumption of the replicator dynamic: revision to the optimal predictor. By comparing the counterfactual forecasts to the actual price in a period, we can identify which counterfactual forecast performed best in a given period. We then can look to see if subjects who changed forecast strategy in the next period chose the same strategy that we identified as the best.

Table 4 reports the results for the first and second announcements across all treatments. The first column reports the number of individual observations for each announcement and treatment. The second column reports the percentage of individual that we observe not changing their strategy between the two periods of interest. The third, fourth, and fifth columns compare the mean of the squared forecast errors for the individuals who we identified as changing their strategy compared to those who we identify as not changing, along with the difference between the two. The final column reports the proportion of participant who we identify as changing their strategy to the strategy that would have been the best to use in the previous period as assumed under the replicator dynamic.

We find evidence consistent with our replicator assumption on all three counts. We find that a significant portion of subjects do not update their strategy following the announcement period. Moreover, the subjects who we do observe changing their strategy experienced significantly larger forecast errors on average than those who do not update in all but one case (second announcement in T3 x A3). We also find that when subjects do change their strategies that they are significantly more likely to change to the strategy that would have performed the best in the previous period as compared to if they were choosing a new strategy at random, again, in all but one case (second announcement in T2 x A3).

Finally, the unified model predicts that we should see increasing depth of reasoning over time converging in the limit to the PFE when $|\beta| < 1$. We can test for whether this condition is predictive by looking at the distribution of strategies that are played across the same subjects in treatments with two announcements. Our propositions predict that over time we

[25] We have tried a number of ϕ values and the results are fairly robust for any relatively large ϕ values. We have even estimated individual specific ϕ values and used those to construct the individual counterfactual. The results do not qualitatively change.

[26] Appendix E provides additional experimental results that details the findings for this classification.

Table 4: Revisions and loss

	Obs	% who do not change strategy	Mean of SFE in p =20/45 No change	Change	Difference	% who change to best strategy
First announcement						
T1 x A2 and A3	84	38.1%	166.9	378.5	211.5	40.4%
					(2.717)	(2.353)
T2 x A2 and A3	90	48.9%	363.8	724.2	360.4	43.5%
					(2.395)	(2.659)
T3 x A2 and A3	120	30.8%	123.3	219.9	96.6	54.2%
					(2.258)	(5.646)
Second announcement						
T1 x A3	42	54.8%	30.8	6867.0	6836.2	78.9%
					(1.596)	(4.988)
T2 x A3	48	45.8%	78.9	595.1	516.2	26.9%
					(3.687)	(0.208)
T3 x A3	66	25.8%	2024.2	1404.4	-619.5	40.8%
					(-0.424)	(2.349)

Notes: This table compares the forecast errors in the period of the announcement for individuals separated by whether we classify them as changing their forecasting strategy in the following period. 'No change' means that we classify them as following the same Level-k strategy in the announcement period and period after. Classifications are done by comparing a subject's forecast to a counterfactual level-k forecast. T-statistics are reported below the 'Difference' results. The 'percent who change to best strategy' are those who we classify as switching to the counterfactual strategy that would have had lowest error in the announcement period. Z-scores for the test of null hypothesis that subjects switched to one of the five level-k strategy at random are reported below.

should see more agents using higher level-k forecasts for the T1 and T3 treatments but not necessarily for T2 treatments. Likewise, our simulations show that T1 and T3 environments generally lead to higher depth of reasoning compared to the T2 treatments.

Figure 16 shows the distribution of forecasts for levels-0 to 3 and REE. We can see for T1 and T3 cases, the distribution shifts to the right. More subjects choose higher level forecasts or are consistent with the REE forecast in the second announcement than in the first. Kolmogorov-Smirnov equality of distribution tests rejects the null of equality at the 5% level for both cases. The T2 case, however, shows a different result. Here, we observe a bifurcation where either agents resort to low levels of reasoning or they jump to the REE.

5.5 Discussion

The experimental evidence provides strong support for hypothesis 1. Large negative feedback results in slower convergence to the REE price relative to when $|\beta| < 1$ or even non-convergence in some cases. In addition, agents appear to update their strategy choices differently in experiments with more than one announcement depending on the expectational

feedback. In the T1 and T3 treatments, people on average choose higher depth of reasoning strategies when faced with a second announced change whereas there is a bifurcation in the T2 treatments with more level-0 and REE strategies played in the second announcement than in the first. This provides further evidence of different behavior when eductive stability or the conditions of Proposition 4 are not satisfied.

We also find strong support for hypothesis 2: agents use level-k reasoning in response to announcements. Depending on the how you classify an agent's forecasts, between 70 to 75% of subjects choose either Level-0, 1, 2, 3, or an REE forecast in an announcement period. We also find some support hypothesis 3 of agents changing strategies over time following the assumption of the replicator dynamic. Specifically, we observe subjects changing their forecasts in response to large forecast errors and changing to level-k forecast strategies that would have performed well in the most recent period. The evidence in this case is weaker though because it relies on our broader level-k classification which is subject to some noise. Therefore, we argue that the unified model provides a reasonable approximation of the behavior but not a one-for-one prediction in most cases.

Finally, we find mixed evidence for the fourth hypothesis regarding revisions of forecasts to higher or lower depth of reasoning depending on the feedback assumption. We do find upward revision of depth of reasoning for the T1 and T3 cases with more low level-k forecasts played when comparing first announcements to second announcements along with quicker convergence. Likewise, with the T2 treatments, we do not find that there is quicker convergence across the two cases, and we observe some reductions in the depth of reasoning. However, we also observe many more subjects playing the REE forecast than predicted by the unified theory in this case. This result suggests that a non-strategic choice of the PFE as a predictor rule may be needed to fully rationalize all the experimental results.

6 Extension: A forward-looking model

A principle motivator for the unified model are the puzzles that arise in forward-looking model to anticipated future events. So far, the model that we have studied only looks one period ahead. However, we find that the intuition from the one-period-ahead case translate quite naturally to the fully forward-looking case.

To illustrate, consider the univariate asset-pricing model in which y_t is the price of the asset and γ_t is the after-tax dividend at time t. Suppose we are initially at $t = 0$ in the steady state and consider a step change in γ beginning at t^* that is that is announced at $t = 0$ (e.g. a decrease in future taxes on dividends). For agents with $k \geq 1$ this information is used. Again in this section we allow for arbitrary K, but set $\omega_{\mathcal{R}} = 0$. Allowing more generally for a known time-varying path γ_t, we consider the model

$$y_t = \gamma_t + \beta \sum_{k=0}^{K} \omega_k E_t^k y_{t+1}. \tag{8}$$

A special case of interest is a single step change $\gamma_t = \gamma$ for $t < t^*$ and $\gamma_t = \gamma'$ for $t \geq t^*$. Level-1 agents form expectations under the assumption that all other agents are level-0, but level 1 agents make use of their knowledge of the path of γ_t. Level k agents use their knowledge of the path of γ_t and assume that the path of y_t is generated by γ_t and level $k-1$ forecasts.[27] Thus level k forecasts at time t satisfy

$$\begin{aligned} E_t^k y_{t+1} &= \gamma_{t+1} + \beta E_{t+1}^{k-1} y_{t+2}, \\ E_{t+1}^{k-1} y_{t+2} &= \gamma_{t+2} + \beta E_{t+2}^{k-2} y_{t+3}, \\ &\vdots \\ E_{t+k-1}^1 y_{t+k} &= \gamma_{t+k} + \beta a, \text{ or} \\ E_t^k y_{t+1} &= \sum\nolimits_{i=1}^{k} \beta^{i-1} \gamma_{t+i} + \beta^k a. \end{aligned}$$

Incorporating these expectations into the model (8) provides the following actual law of motion

$$y_t = \gamma_t + \sum_{i=1}^{K} \sum_{j=1}^{i} \omega_i \beta^j \gamma_{t+j} + \left(\beta \sum_{i=0}^{K} \omega_i \beta^i \right) a.$$

Under adaptive learning we assume as before that

$$E_{t+j}^0 y_{t+j+1} = a_{t-1} \text{ for all } j \geq 0$$

where a_t evolves according to (3).[28] We observe that if γ_t is constant then this ALM reduces to (2) with $\omega_{\mathcal{R}} = 0$ and finite K, i.e.

$$y_t = \gamma \left(1 + \frac{\beta}{1-\beta} \sum_{k=1}^{K} \left(1 - \beta^k\right) \omega_k \right) + \left(\beta \sum_{k=0}^{K} \beta^k \omega_k \right) a_{t-1}.$$

This implies that the stability results of Section 3.1 continue to hold.

Remark 1 *Level k agents are assumed to obey the "anticipated utility" assumption that they act as if a_{t-1} will not change in future periods. This assumption, which is standard in the adaptive learning literature, could of course be modified if agents are thought to have even greater sophistication.*

6.1 Anticipated events

The unified model in this case implies a natural interpretation to forward-looking behavior. The number of periods an agent looks into the future when making a decision is directly

[27] We note that this assumption implies that level k agents believe all other agents are level $k-1$ agents, and further, level k agents believe that all level $k-1$ agents believe that all other agents are level $k-2$ agents, and so on.

[28] We are assuming that agents update a_t before y_t is realized.

linked to her level of k. In fact, a level-k agent looks exactly k periods ahead. Therefore, even in the case where an announced policy change is fully credible and understood by all market participants, there may not be a contemporaneous response because everyone is using a strategy whose horizon is too short to incorporate the impact. In this way, the unified dynamic rationalizes finite horizon approaches such as that proposed by Branch, Evans and McGough (2013) and Woodford (2019).

Figure 17 shows a simulation of announced change to γ in the unified model with both adaptive learning and the replicator dynamic active. The optimal k actually drops as the structural change nears. The agents anticipate that their less sophisticated counterparts will not take the announcement into account. Therefore, they do not respond to the impending change. This further weakens the effect of the anticipated change.

7 Conclusion

The union of behavioral heterogeneity, adaptive learning, and level-k reasoning brings together three behavioral assumptions that enjoy wide experimental support. Level-k reasoning in particular is found to be a good description of how people form higher order beliefs in wide variety of settings. We contribute to this literature by showing how level-k beliefs naturally fit with some of the most common forms of bounded rationality studied in macroeconomic environments. In addition, we provide a plausible way in which level-k beliefs may evolve over time in response to forecast errors and in response to learning through the level-0 forecast.

Our experiment provides compelling evidence for the key features of the unified model. We observe heterogeneous behavior that is consistent with level-k deductions as well as changes in behavior over time that are consistent with revisions to participants' depth of reasoning following the assumptions of our replicator dynamic. These results are important because it shows that insights from beauty contest experiments can extend to more complicated dynamic settings and because it provides experimental support for using these mechanisms to model boundedly rational response to anticipated events. The latter of which has broad applications in macroeconomics.

References

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

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

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

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

Branch, William A, "The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations," *The Economic Journal*, 2004, *114* (497), 592–621.

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

Branch, William, George W Evans, and Bruce McGough, "Finite horizon learning," *Macroeconomics at the Service of Public Policy*, 2013, pp. 141–163.

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

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

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

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

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

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

Evans, George W, "Expectations in macroeconomics adaptive versus eductive learning," *Revue économique*, 2001, *52* (3), 573–582.

____ and Bruce McGough, "Interest-Rate Pegs in New Keynesian Models," *Journal of Money, Credit and Banking*, 2018, *50* (5), 939–965.

____ and Garey Ramey, "Expectation calculation and macroeconomic dynamics," *The American Economic Review*, 1992, pp. 207–224.

____ and ____, "Calculation, adaptation and rational expectations," *Macroeconomic Dynamics*, 1998, *2* (2), 156–182.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Negro, Marco Del, Marc P Giannoni, and Christina Patterson, "The forward guidance puzzle," *FRB of New York Staff Report*, 2012, (574).

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

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

Woodford, Michael, "Monetary policy analysis when planning horizons are finite," *NBER Macroeconomics Annual*, 2019, *33* (1), 1–50.

Appendices

A Proofs

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

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

Proof. Write $T_\gamma(a, \omega, \beta)$ as the realized value of y given the datum $(\beta, \omega, a, \gamma)$. Then

$$\begin{aligned} T_\gamma(a, \omega, \beta) - \frac{\gamma}{1-\beta} &= \gamma + \frac{\beta\gamma}{1-\beta}\sum_{k\geq 0}\omega_k - \frac{\beta\gamma}{1-\beta}\sum_{k\geq 0}\beta^k\omega_k + \beta a\sum_{k\geq 0}\beta^k\omega_k - \frac{\gamma}{1-\beta} \\ &= \frac{\gamma}{1-\beta} - \left(\frac{\gamma}{1-\beta}\right)\beta\sum_{k\geq 0}\beta^k\omega_k + \beta a\sum_{k\geq 0}\beta^k\omega_k - \frac{\gamma}{1-\beta} \\ &= T_0\left(a - \frac{\gamma}{1-\beta}, \omega, \beta\right). \end{aligned} \tag{9}$$

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

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

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

Now observe that

$$\begin{aligned} \arg\min_{k\in\mathbb{N}} \phi^\varepsilon(\beta, k, a, 0) &= \arg\min_{k\in\mathbb{N}} \left| a\beta^k - a\sum_n \beta^n\omega_n \right| = \arg\min_{k\in\mathbb{N}} |a| \left| \beta^k - \sum_n \beta^n\omega_n \right| \\ &= \arg\min_{k\in\mathbb{N}} \left| \beta^k - \sum_n \beta^n\omega_n \right| = \arg\min_{k\in\mathbb{N}} \phi^\varepsilon(\beta, k, 1, 0). \end{aligned} \tag{10}$$

Also, by (9) we have that

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

so that

$$\arg\min_{k\in\mathbb{N}} \phi^\varepsilon(\beta, k, a, \gamma) = \arg\min_{k\in\mathbb{N}} \phi^\varepsilon(\beta, k, a - \bar{y}, 0). \tag{11}$$

Putting (10) and (11) together yields

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

which completes the proof. ■

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

$$T_\gamma(a,\omega) = \gamma\left(1 + \frac{\beta}{1-\beta}\sum_{k\geq 0}(1-\beta^k)\omega_k\right) + \beta\sum_{k\geq 0}\beta^k\omega_k a \tag{12}$$

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

$$a_t = a_{t-1} + \phi\left(T_\gamma\left(a_{t-1},\omega_{t-1}\right) - a_{t-1}\right). \tag{13}$$

We have the following result.

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

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

Proof. First, observe that (9) and (13) imply

$$\begin{aligned} a_t - \frac{\gamma}{1-\beta} &= a_{t-1} - \frac{\gamma}{1-\beta} + \phi\left(T_\gamma\left(a_{t-1},\omega_{t-1}\right) - \frac{\gamma}{1-\beta} - \left(a_{t-1} - \frac{\gamma}{1-\beta}\right)\right) \\ &= a_{t-1} - \frac{\gamma}{1-\beta} + \phi\left(T_0\left(a - \frac{\gamma}{1-\beta},\omega_{t-1}\right) - \left(a_{t-1} - \frac{\gamma}{1-\beta}\right)\right), \end{aligned}$$

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

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

$$\left|\beta\sum_{k\geq 0}\beta^k\omega_k\right| \leq |\beta|\sum_{k\geq 0}|\beta^k|\omega_k \leq |\beta|\sum_{k\geq 0}|\beta|\,\omega_k \leq \beta^2. \tag{14}$$

Next, write the recursion (13) as

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

By equation (14),

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

It follows that

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

establishing item 1.

Now let $\beta > 1$. The same reasoning as in (15), but with the inequalities reversed, yields

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

It follows that

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

and the result follows. ■

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

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

Lemma A.3 *Suppose $|\beta| < 1$ and $\gamma = 0$. Then $k < \hat{k}(y)$ implies that there exists $\delta \in (0,1)$ such that $|y| < (1-\delta)|a\beta^k|$.*

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

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

Thus we compute

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

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

Proof. Let $t \geq N$. Since $\hat{k}(y_t) > \tilde{k}$, it follows from Lemma A.3 that $(1-\delta)|a\beta^{\tilde{k}}| > |y_t|$, for some $\delta \in (0,1)$. Thus $n \leq \tilde{k}$ implies $|a\beta^n - y_t| \geq |a\beta^n| - |y_t| > |a\beta^n| - (1-\delta)|a\beta^{\tilde{k}}| > 0$. Using this estimate in the replicator yields, for $s \geq 1$,

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

Because $r(0) \geq 0$ and $r' > 0$, it follows that $\lim_{t\to\infty} \omega_{nt+s} = 0$ as $s \to \infty$. ■

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

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

We are now ready to prove the main result.

Proof of Theorem 2. By Lemma A.1 we may assume $\gamma = 0$. To simplify notation, let $\hat{k}_t = \hat{k}(y_t)$. Noting that if all the weights are driven to zero then $y_t \to 0$, we see that by Lemma A.4 it suffices to show $\hat{k}_t \to \infty$. It is helpful to introduction some notation: for $y \in \mathbb{R}$ and $m(y), n(y) \in \mathbb{N}$, write $m(y) \succ n(y)$ when the level-m forecast is superior to the level-n forecast, i.e.,

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

Now set $\tilde{k} = \liminf \hat{k}_t$, and let N be chosen as in Lemma A.5 . To show $\hat{k}_t \to \infty$ it suffices to show that $\tilde{k} = \infty$. Our proof strategy is to assume $\tilde{k} < \infty$ and derive a contraction. To this end, it suffices to find some $M > 0$ so that $t \geq M$ implies the existence of $m(y_t) > \tilde{k}$ with $m(y_t) \succ \tilde{k}$, as this contradicts the definition of $\tilde{k}$ as the limit infimum of the $\hat{k}_t$. The remainder of the proof, which is developed in three steps, derives this contradiction.

Step 1. We establish the following claim:

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

<u>Proof of claim.</u> We know that for all $t \geq N$ we have $\hat{k}_t \geq \tilde{k}$. For $n \geq 1$ let $\Omega_t^l(n) = \sum_{k<n} \omega_{kt}$, and note that since $|y_t| \leq |a\beta|$ it follows from $\tilde{k} \geq 1$. Then

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$$\begin{aligned} 2|\beta|^{\tilde{k}+1}(1+\varepsilon) &< |\beta|^{\tilde{k}+2}+|\beta|^{\tilde{k}}, \text{ or} \\ 0<|\beta|^{\tilde{k}+1}(1+\varepsilon)-|\beta|^{\tilde{k}+2} &< |\beta|^{\tilde{k}}-|\beta|^{\tilde{k}+1}(1+\varepsilon). \end{aligned} \tag{15}$$

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

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

Case 2: $\hat{k}_t=\tilde{k} \bmod 2$. If $y_t<0$ then $a\beta^{\tilde{k}}$ is negative. Thus

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

where the first inequality follows from (15). Thus

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

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

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

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

where, again, the left inequality follows from (15). Thus

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

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

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

The dynamic system for a_t may be written

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

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

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

This completes the proof of items 3(a) and 3(b).

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

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

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

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

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

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

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

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

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

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

Now let

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

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

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

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

Since ψ_β is continuous, it follows that $\Omega_0^\pm$ are closed. Since no-where denseness is closed under finite unions, it suffices to show that the interiors of $\Omega_0^\pm$ are empty. Thus let $\omega \in \Omega_0^+$. Let $N \in \mathbb{N}$ so that $\omega \in \Delta^N$. Since $\beta < -1$ and $\psi_\beta(\omega) = 1$ there is an even $n \in \mathbb{N}$ and an odd $m \in \mathbb{N}$, with $n, m \leq N$ and such that $\omega_n, \omega_m \neq 0$. For $k \in \mathbb{N}$ with $k \geq 2$, define $\omega^k \in \Delta^N \subset \Omega$ as follows:

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

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

Characterization of $\hat{k}$: Proposition 1. Some preliminary work is required. By Lemma A.1 we may assume $\gamma = 0$ and $a = 1$. Next, recall from Section 3.2 for $K \in \mathbb{N}$ we define $\omega^K = \{\omega_n\}_{n=0}^K$ be a weight system with weights given as $\omega_n = (K+1)^{-1}$. It follows that

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

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

It is helpful to define k^* as the *continuous* counterpart to $\hat{k}$. Care must be taken to accommodate $\beta < 0$. We define k^* as follows:

$$k^*\left(K, \beta\right) = \frac{\log\left(\psi\left(K, \beta\right)^2\right)}{\log(\beta^2)}. \tag{16}$$

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

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

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

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

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

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

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

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

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

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

This function has the following properties:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Case 4: $\beta < -1$ and $\psi > 0$. Then $k^* \geq 0$ (by assumption) and $\hat{k}$ is even. Also $\alpha > 0$ implies $\phi_\alpha(\alpha, \beta) > 0$. Thus if $\lfloor k^* \rfloor_{\text{even}} > 2$ and if k_i are even with $k_1 < \lfloor k^* \rfloor_{\text{even}}$ and $k_2 > \lceil k^* \rceil_{\text{even}}$, then

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

It follows that near and below $\beta = 1$ we have $\frac{K}{2}, \frac{K+4}{2} \prec \frac{K+2}{2}$.

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

$$\Delta(1, 2, \beta) = (2 - 4(1+K)^{-1})(\beta-1)^2 + \mathcal{O}\left(|\beta-1|^3\right), \tag{17}$$

so that $\hat{k} = 2$ for $K \geq 2$. When $K = 1$ we again find $\Delta = 0$, so that $\hat{k} = 1$.

Case 2: $-1 < \beta < 0$. We establish item 2(a) by direct analysis, and noting that it suffices to study the behavior of k^*. Noting that $-1 < \psi < 0$, we compute

$$\log \psi^2 = 2\log(-\psi) = \log\left(\frac{\beta}{\beta-1}\right) + \log\left(1 - \beta^{K+1}\right) - \log(1+K) \to -\infty \tag{18}$$

$$K^{-1}\log \psi^2 = K^{-1}\log\left(\frac{\beta}{\beta-1}\right) + K^{-1}\log\left(1 - \beta^{K+1}\right) - K^{-1}\log(1+K) \to 0 \tag{19}$$

Since $k^* = \log \psi^2 / \log \beta^2$ and $\log \beta^2 < 0$ we see that by equation (18) $k^* \to \infty$, and that by equation (19) ${}^{k^*}\!/\!_K \to 0$.

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

have $\hat{k} \in \{1, 2\}$. The expansion (17) then shows that $\hat{k} = 3$ for $K \geq 2$. Also as before, $K = 1$ implies $\Delta = 0$, so that $\hat{k} = 1$. Finally, for item 2(c), we find using we find that $\beta \to -1^+$ implies $k^* \to \infty$, and the result follows.

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

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

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

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

It follows that

$$\lim_{K\to\infty} \frac{k^*}{K} = \lim_{K\to\infty} (K\log(-\beta))^{-1}\log\left((\beta^2)^{\frac{K+1}{2}} + (-1)^{K+1}\right). \tag{20}$$

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

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

It follows that

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

which, when combined with (20), yields the result.

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

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

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

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

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

Appendix D: Experiment materials

Negative feedback case

Computer based tutorial:

- What is your role?

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

- What makes you an expert in this market?

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

- What is a widget?

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

- Why do the firms need to forecast the price?

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

- How am I paid?

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

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

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

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

- The Demand for Widgets:

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

  $$q = A - Bp$$

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

- The Supply of Widgets:

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

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

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

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

- Prices and Expected Prices:

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

- Why does this occur?

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

  By equating supply and demand,

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

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

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

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

- A bit of randomness:

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

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

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

Positive feedback case

Computer based tutorial:

- What is your role?

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

- What makes you an expert in this market?

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

- What is a widget?

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

- Why do the firms need to forecast the price?

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

- How am I paid?

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

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

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

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

- The Demand for Widgets:

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

  $$q = A + Bp$$

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

- The Supply of Widgets:

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

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

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

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

- Prices and Expected Prices:

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

- Why does this occur?

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

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

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

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

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

- A bit of randomness:

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

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

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

Paper instructions:

Exit Survey

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

    a. The history of market prices

    b. The market equations

    c. The history of my own price forecasts

    d. The history of my own forecasts errors

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

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

    a. The history of market prices

    b. The market equations

    c. The history of my own price forecasts

    d. The history of my own forecasts errors

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

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

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

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

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

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

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

    a. The history of market prices

    b. The market equations

c. The history of their own price forecasts

d. The history of their own forecasts errors

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

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

   a. The history of market prices

   b. The market equations

   c. The history of their own price forecasts

   d. The history of their own forecasts errors

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

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

Appendix E: Additional Experimental Results

Table 5: Classification of predictions using counterfactual forecast rules

	All observations			Announcement periods only		
Treatment	RMSE	RMedSE	70th Pctl	RMSE	RMedSE	70th Pctl
T1 x A1	15.16	0.23	0.43	73.14	0.54	1.21
T1 x A2	10.67	0.26	0.58	4.40	1.23	2.80
T1 x A3	9.99	0.31	0.72	30.39	1.37	2.76
T2 x A1	10.18	0.25	0.68	3.20	0.55	3.00
T2 x A2	17.83	0.30	1.09	4.53	1.08	3.80
T2 x A3	4.20	0.50	1.32	5.46	2.00	5.00
T3 x A2	23.95	0.40	1.00	14.32	1.27	7.57
T3 x A3	16.57	0.34	1.00	33.68	1.61	4.86

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

Figure 15: Evidence of level-k reasoning in announcement periods

Announcement 1 | Announcement 2

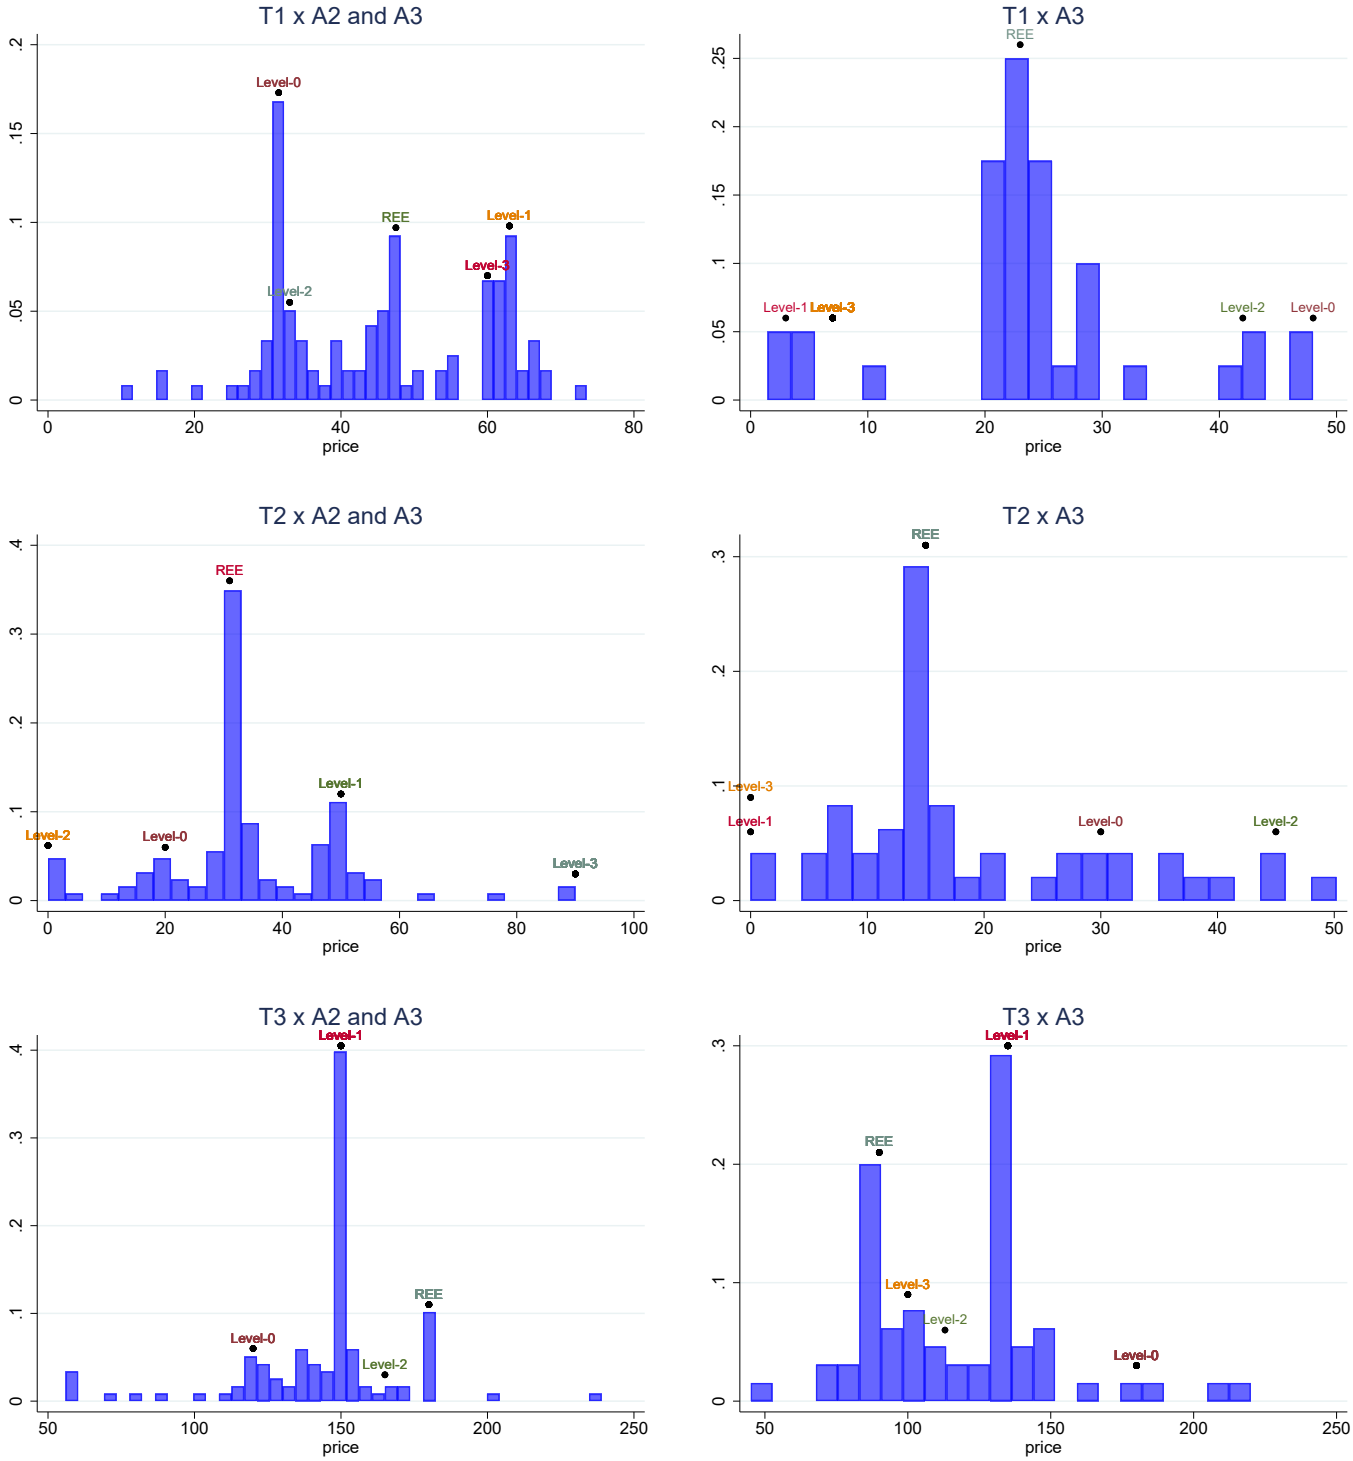

Notes: Histogram of predictions after an announced change. Labels correspond to the level-0, level-1, level-2, and level-3 forecasts implied by a level-0 beliefs of the price at the REE value before the announced change. The REE label corresponds to the new REE price implied by the announced change.

Figure 16: Distribution of strategies over successive announcements

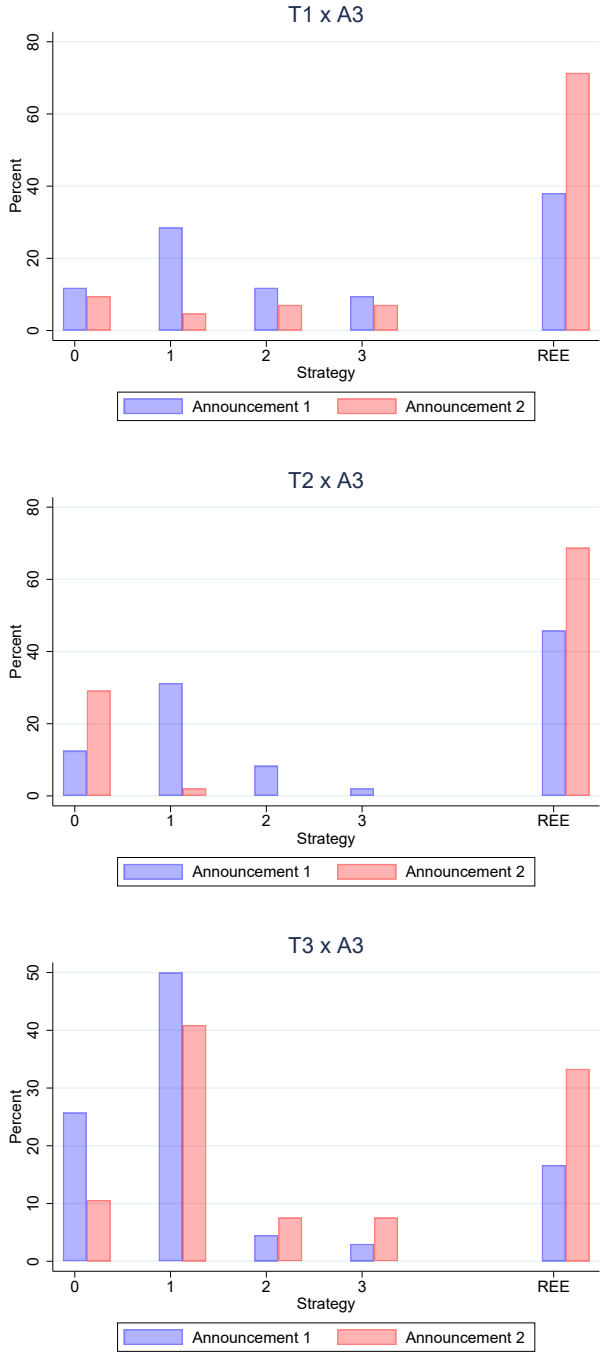

Notes: Classification of the depth of reasoning is done by forecast errors from counterfactual level-k forecasts to the observed forecasts.

Figure 17: Forward-looking model with Unified Dynamics

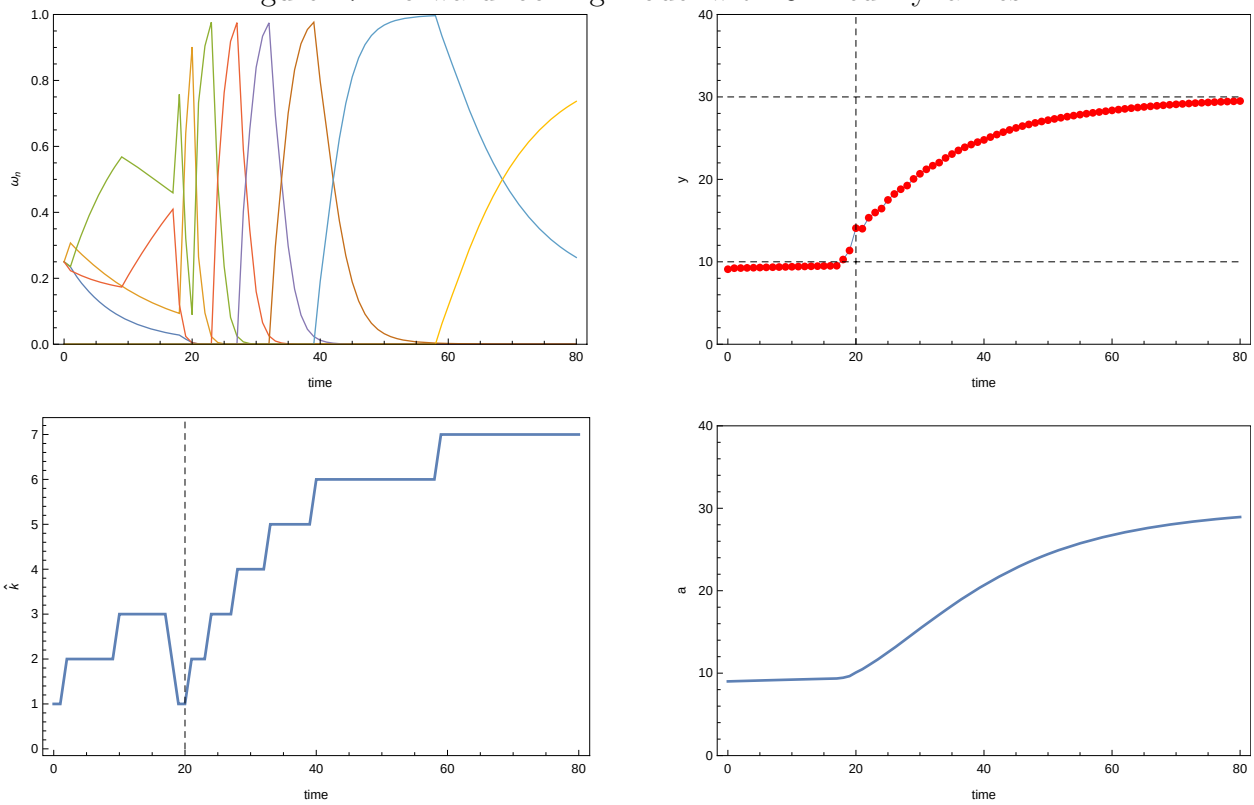

Notes: Simulation of announced change of γ from 1 to 2 in period 20. The shift is known to all agents in period 1. The remaining parameters are $\beta = 0.9$, $\phi = 0.1$, and $\alpha = 1$. The simulation start with agents equally distributed across levels 0, 1, 2, and 3.

Widget Game Instruction Summary:

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

Widget Game Rules

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

Widget Game Instruction Summary:

- Your job is to forecast the price of a widget next period
- Demand for widgets is determined by the market price
  - $q = A + B\,p$
- The total supply of widgets to the market is determined by the **average** of all price forecasts submitted to the market
  - $q = C + D$ x average price forecast
- Combining supply and demand, we have the **key formula** that determines price in the market
  - $P = (C - A)/B + D/B$ x average expected price + *noise*
    - Recall that *noise* is small and **on average equal to zero**
- **An Example:** $A = 0,\ B = 2,\ C = 60,\ D = 1,$ and $noise = 0$, what is price if the average price forecast is 42?
  - $p = 30 + ½$ x average price forecast
  - $P = 30 + ½ \times 42 = 30 + 21 = 51$
- You are paid based on **accuracy of your forecast** according to the following formula
  - Payment = 0.50 – 0.03 $(p - your\ price\ forecast)^2$
    - A perfect forecast in a round earns 50 cents
    - A very poor forecast results in 0.00
- **KEY POINT: The market has positive feedback. Therefore, if the average price forecast is high, the market price will be high. And, if the average price forecast is low, then the market price will be low.**
- **Your Notes:**
  - **–**
  - **–**
  - **–**
  - **–**
  - **–**
  - **–**

Widget Game Rules

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

Figure 18: Exit survey question 3 responses

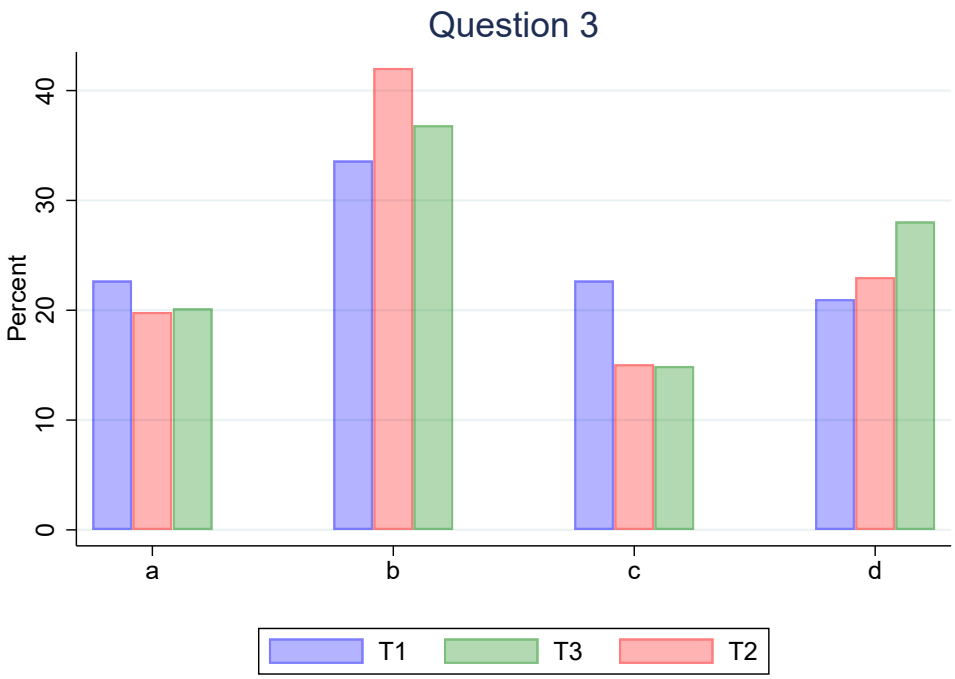

Notes: This figure shows the response to question 3 from exit survey separated by treatment type.

Figure 19: Exit survey results summary

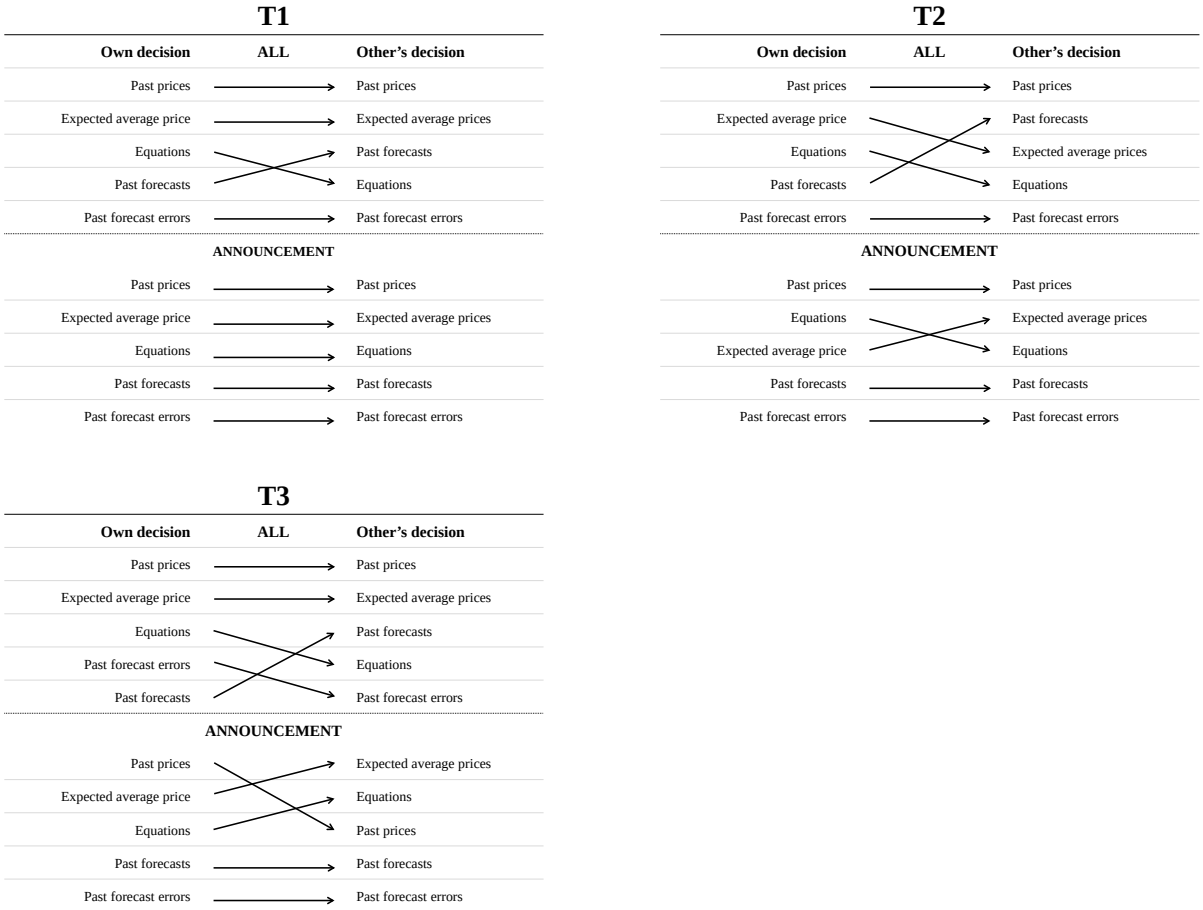

Notes: This figure shows the importance ranking of information to participants that was used to make their "own decision" for the price forecast compared with the information the participant's believed "other's" used when making their decision. The top list of each panel shows the responses for the question regarding forecasts in general while the bottom list shows the responses for decisions specifically asking about announcement periods.