

A Unified Model of Learning to Forecast

Online Appendix

George W. Evans

University of Oregon and University of St Andrews

Christopher G. Gibbs

Bruce McGough

University of Sydney

University of Oregon

December 30, 2023

A1 PROOFS

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

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

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

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

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

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

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

Now observe that

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

Also, by (A1) we have that

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

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

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

Putting (A2) and (A3) together yields

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

which completes the proof. ■

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

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

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

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

We have the following result.

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

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

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

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

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

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

$$\left| \beta \sum_{k \geq 0} \beta^k \omega_k \right| \leq |\beta| \sum_{k \geq 0} |\beta^k| \omega_k \leq |\beta| \sum_{k \geq 0} |\beta| \omega_k \leq \beta^2. \quad (\text{A6})$$

Next, write the recursion (A5) as

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

By equation (A6),

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

It follows that

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

establishing item 1.

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

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

It follows that

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

and the result follows. ■

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

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

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

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

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

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

Thus we compute

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

We are now ready to prove the main result.

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

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

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

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

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

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

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

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

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

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

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

Step 1. We establish the following claim:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

where the first inequality follows from (A7). Thus

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

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

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

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

where the reasoning is as above. Thus

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

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

Proof of Theorem 3. Lemma A.2 establishes items 1 and 2, and so we focus here only on item 3. Also, as demonstrated in the proof of Lemma A.2, we may assume $\gamma = 0$. We recall the notation $\Omega = \dot{\cup}_n \Delta^n$ and $\psi_\beta : \Omega \rightarrow \mathbb{R}$, 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 $|A(\beta, \omega, \phi)| < 1 \implies a_t \rightarrow 0$ and $|A(\beta, \omega, \phi)| > 1 \implies |a_t| \rightarrow \infty$. We compute

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

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

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

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

Assuming $\psi_\beta^n : \Delta^n \rightarrow \mathbb{R}$ is continuous, we have that $(\psi_\beta^n)^{-1}(U)$ 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 $\alpha \rightarrow \omega_\alpha$ continuously maps $(0, 1)$ into Δ^{2n+1} , whence into Ω . Let $\Psi_\beta : (0, 1) \rightarrow \mathbb{R}$ 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

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

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 $y_t \rightarrow \bar{y}$ and $\omega \in \Omega_u$ implies $|y_t| \rightarrow \infty$. 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 Ω_0^\pm are closed. Since no-where denseness is closed under finite unions, it suffices to show that the interiors of Ω_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 $\omega^k \rightarrow \omega$. Owing to the construction of the direct-limit topology, we have that $\omega^k \rightarrow \omega$ in Ω as well. Thus, given an arbitrary element $\omega \in \Omega_0^+$ we have constructed a sequence in Ω_u converging to it, and since $\Omega_u \cap \Omega_0^+$ is empty, we conclude that ω is not in the interior of Ω_0^+ . So the interior of Ω_0^+ is empty, and since Ω_0^+ is closed, we conclude that Ω_0^+ is nowhere dense. The same argument applies to Ω_0^- , which shows that $\Omega_0 = \Omega_0^- \dot{\cup} \Omega_0^+$ is no-where dense. ■

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

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

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

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

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

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

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

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

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

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

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

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

3. Suppose $\beta < -1$

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

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

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

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

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

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

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

$$k^*(K, \beta) = \frac{\log(\psi(K, \beta)^2)}{\log(\beta^2)}. \quad (\text{A9})$$

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

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

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

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 $\hat{k} \in \{k_{\text{low}}^*, k_{\text{high}}^*\}$.*

Proof. We begin with the following observations on the parity of \hat{k} .¹ Recall that 0 is taken as even.

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

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

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

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

This function has the following properties:

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

¹The *parity* of $n \in \mathbb{N}$ is its equivalence class mod 2. Thus n and m have the same parity if they are either both even or both odd.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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 $K \rightarrow \infty$ implies $k^* \rightarrow \infty$ and $k^*/K \rightarrow 0$.

```
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 $\beta \rightarrow 1^-$ implies $k^* \rightarrow K/2 + 1$. 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 $\beta = 1$,

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

so that when $K \geq 3$ and β is near and below 1, we conclude that $\Delta < 0$, so that $\hat{k} = 1/2(K+1)$. When $K = 1$ a direct computation shows $\Delta = 0$, so that both $\lfloor k^* \rfloor$ and $\lceil k^* \rceil$ yield the same forecast. Our tiebreaker, then, chooses $\hat{k} = 1$.

Now suppose K is even. Then for β near and below 1 we know that k^* is near $K/2 + 1 \in \mathbb{N}$. Unfortunately, we do not know if k^* approaches its limit monotonically. Thus we can only conclude that for β near and below 1 we have

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

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

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

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

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

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

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

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

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

Since $k^* = \log \psi^2 / \log \beta^2$ and $\log \beta^2 < 0$ we see that by equation (A11) $k^* \rightarrow \infty$, and that by equation (A12) $k^*/K \rightarrow 0$.

Turning to item 2(b), using Mathematica we find that $\beta \rightarrow 0^-$ implies $k^* \rightarrow 1$, and since $\beta \in (0, 1)$, we know that $\psi < 0$ so that \hat{k} is odd. It follows that for β is near and below 0 we have $\hat{k} \in \{1, 2\}$. The expansion (A10) then shows that $\hat{k} = 3$ for $K \geq 2$. Also as before, $K = 1$ implies $\Delta = 0$, so that $\hat{k} = 1$. Finally, for item 2(c), we find using we find that $\beta \rightarrow -1^+$ implies $k^* \rightarrow \infty$, 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 $\alpha \rightarrow \infty$ for $x > 1$ and for any $\beta \in \mathbb{R}$, which shows that $|\psi(K, \beta)| \rightarrow \infty$ as $K \rightarrow \infty$. It follows that $\log \psi^2 \rightarrow \infty$, and thus k^* and \hat{k} go to infinity as $K \rightarrow \infty$. Next note

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

It follows that

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

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

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

It follows that

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

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

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

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

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

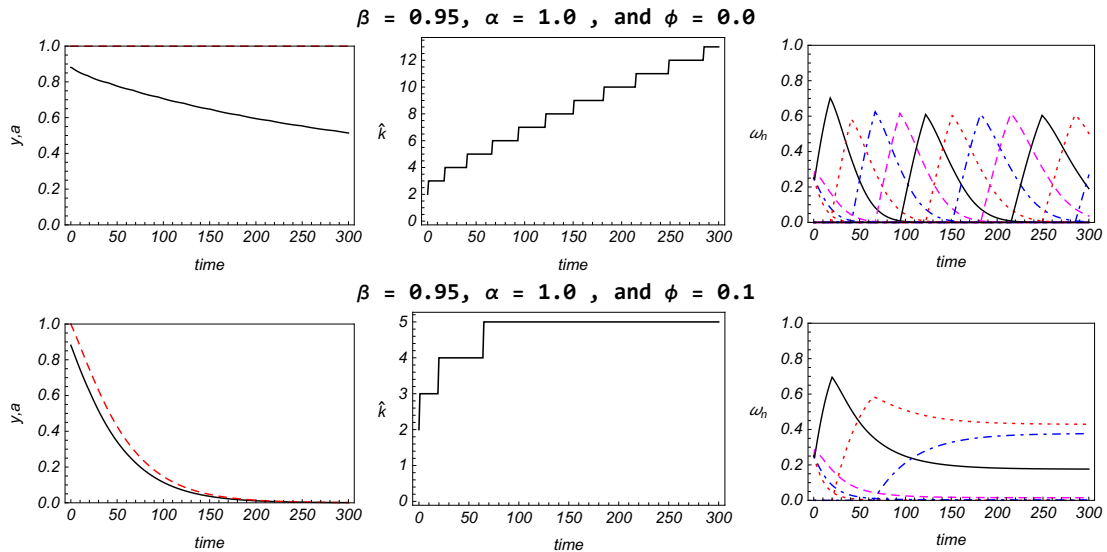
$$\lim_{\beta \rightarrow -\infty} \Delta(K-1, K+1) = \lim_{\beta \rightarrow -\infty} \Delta(K+1, K+3) - \infty,$$

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

A2 SIMULATED DYNAMICS OF THE UNIFIED MODEL

To illustrate how convergence is achieved under different specifications of the unified dynamics, we consider a variety of special cases operating under a range of feedback parameters β . In this section, without loss of generality, we set γ at zero, so that $\bar{y} = 0$ (equivalently, the dynamics for y and a may be viewed as in deviation form). We take the parametric form of the rate function for the replicator dynamics to be given by $r(x) = 2/\pi \tan^{-1}(\alpha x)$, with $\alpha > 0$. Finally, all simulations are initialized with $a_0 = 1$ and $\omega_{k0} = 1/4$ for $k = 0, 1, 2, 3$.

Figure A1: Simulated dynamics with positive feedback



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


We start with with the stable positive feedback case $0 < \beta < 1$: see Figure A1, where $\beta = 0.95$ and $\alpha = 1$. Upper row corresponds to replicator dynamics

($\phi = 0$) and bottom row to unified dynamics ($\phi = 0.1$): we omit results associated with adaptive dynamics as they simply show monotonic convergence of a and y to \bar{y} .

Under replicator dynamics, y exhibits monotone convergence to \bar{y} , as the weight distribution shifts to higher k -level forecasts. The upper-right panel provides the dynamics of agents' weights. The time paths for weights $\omega_{n0} = 1/4$, $n = 0, 1, 2, 3$, are distinguished by plot-style: *red dotted*, *blue dash-dot*, *dashed magenta* and *solid black*, respectively. As the replicator adds higher forecast levels, the associated paths are graphically identified in an analogous fashion by repeating the styles mod four. Under replicator dynamics, lower-level forecasts gradually fall out of favor and are replaced by higher-level forecasts.

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

Figure A2: Simulated dynamics with negative feedback.



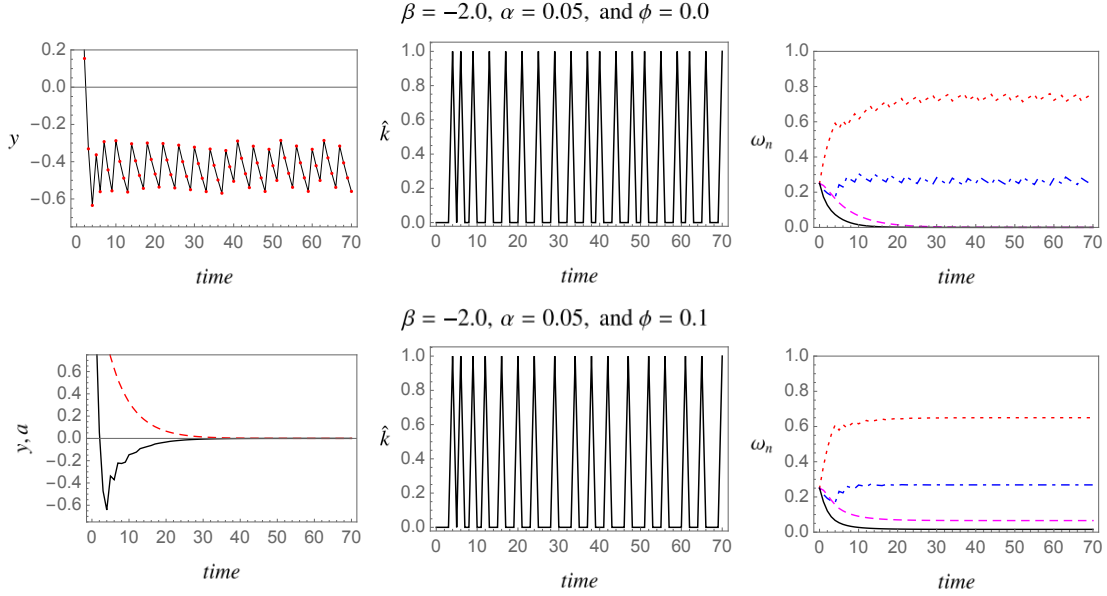
Notes: Simulation of replicator dynamics only (top) and unified dynamics (bottom). $\omega_{n0} = 1/4$ for $n = 0, 1, 2, 3$, and the time paths for these four weights are distinguished by plot-style: *red dotted*, *blue dash-dot*, *dashed magenta* and *solid black*, respectively.

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

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

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

Figure A3: Simulated dynamics with large negative feedback.



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

are fully rational and have common knowledge of the structure they are unable to coordinate on the REE. However, as indicated by Corollary 1, when $\beta < -1$ the REE is stable under adaptive dynamics provided the gain is sufficiently small.

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

A3 SUMMARY STATISTICS AND ADDITIONAL EXPERIMENTAL RESULTS

Table A1 reports summary statistics for the experiment. In total, 372 individuals participated in 62 experimental markets. All T1 and T2 treatments were conducted in May and June of 2018 at the UNSW Sydney BizLab. Two sessions for each treatment were scheduled with the aim of testing at most five markets in each session. Participant no-shows account for the different number of markets across the treatments.³ 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.⁴ Again, no-shows account

²We find numerically that there are at least two stable 11-cycles.

³Subjects were recruited using ORSEE (see Greiner, 2015).

⁴The different session sizes at the University of Sydney reflect lab capacity constraints due to equipment issues and subject recruitment limitations. Christopher Gibbs left UNSW for Sydney in July of 2018, causing a the delay between experiments.

Table A1: Summary statistics

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

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

for the different number of experimental markets across treatments.

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

We observed more outliers in individual predictions in this study than were observed, for example, in Bao and Duffy (2016). However, we also have more than double the participants. Some outliers are easily explained as “fat finger” errors where an extra zero is added to a forecast. Others reflect participants with a penchant for anarchy who consistently typed in nonsensical forecasts. In fact, we identify two anarchists who repeatedly typed in the highest price permitted just to see what would happen. One of these anarchists actually provided a nice natural experiment within our laboratory experiment, which we discuss in detail in Section A4 below.

When classifying individual forecasts without cutoffs, we chose to not classify 35 out of the 18,367 forecasts from our analysis (5 of which occurred in announcement rounds out of 517 observations in total).⁵ Nearly half of the total outliers forecasts were submitted by just 3 (out of the 372) participants in the study. The outlier predictions on average were for a price of 391, which is nearly 200 larger than any plausible price in any treatment. If these outliers were classified as level- k , then most are classified as an REE prediction (e.g. in a positive feedback treatment when level- k forecasts converge from below the REE price and the outlier is above the REE price) or a level-0 prediction (e.g. when convergence starts from above an REE price and level- k deductions are closer to the REE price), which is clearly not in keeping with what the classification is attempting to achieve.⁶

⁵The odd number of observations is due to two markets that did not complete the experiment. One market in a T2×A1 treatment ended early when a participant withdrew from the experiment. The other was a T2×A2 treatment that ended a short time after the announcement when a student kicked a power cord knocking out two computers with players in the same market. The data up to that point was saved, but there was no way to let the students pick up where they left off.

⁶Inclusion of these outliers actually makes some of our results stronger. For example, with respect to the result reported in Table A4, the forecast errors generated by some of these outliers

Figure A4: Average market price relative to REE

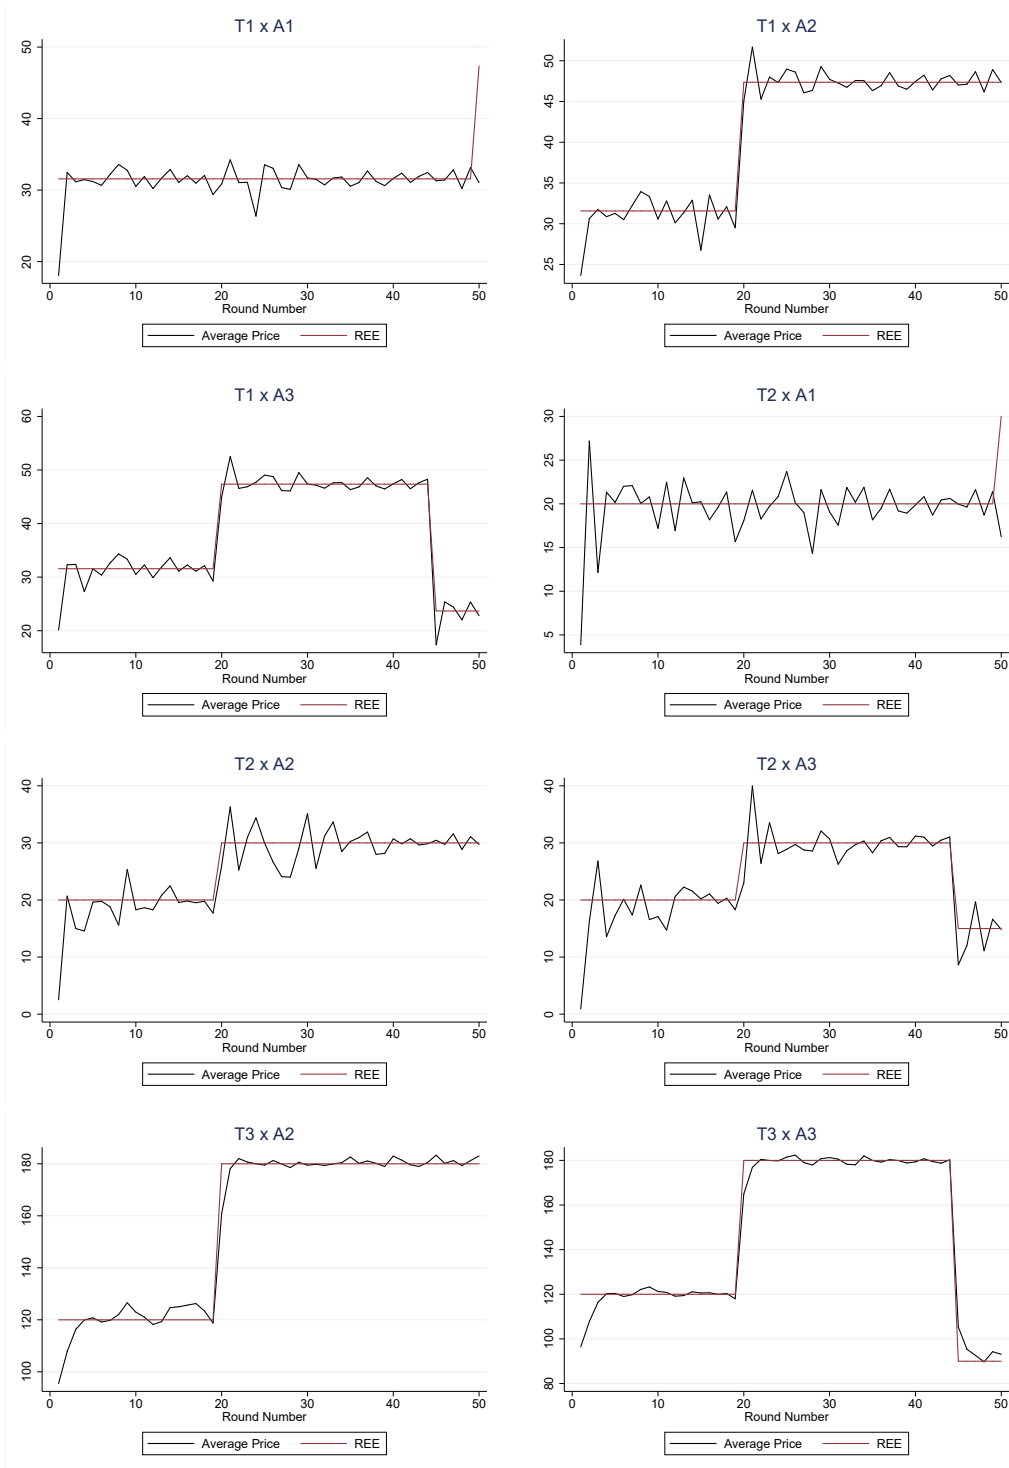
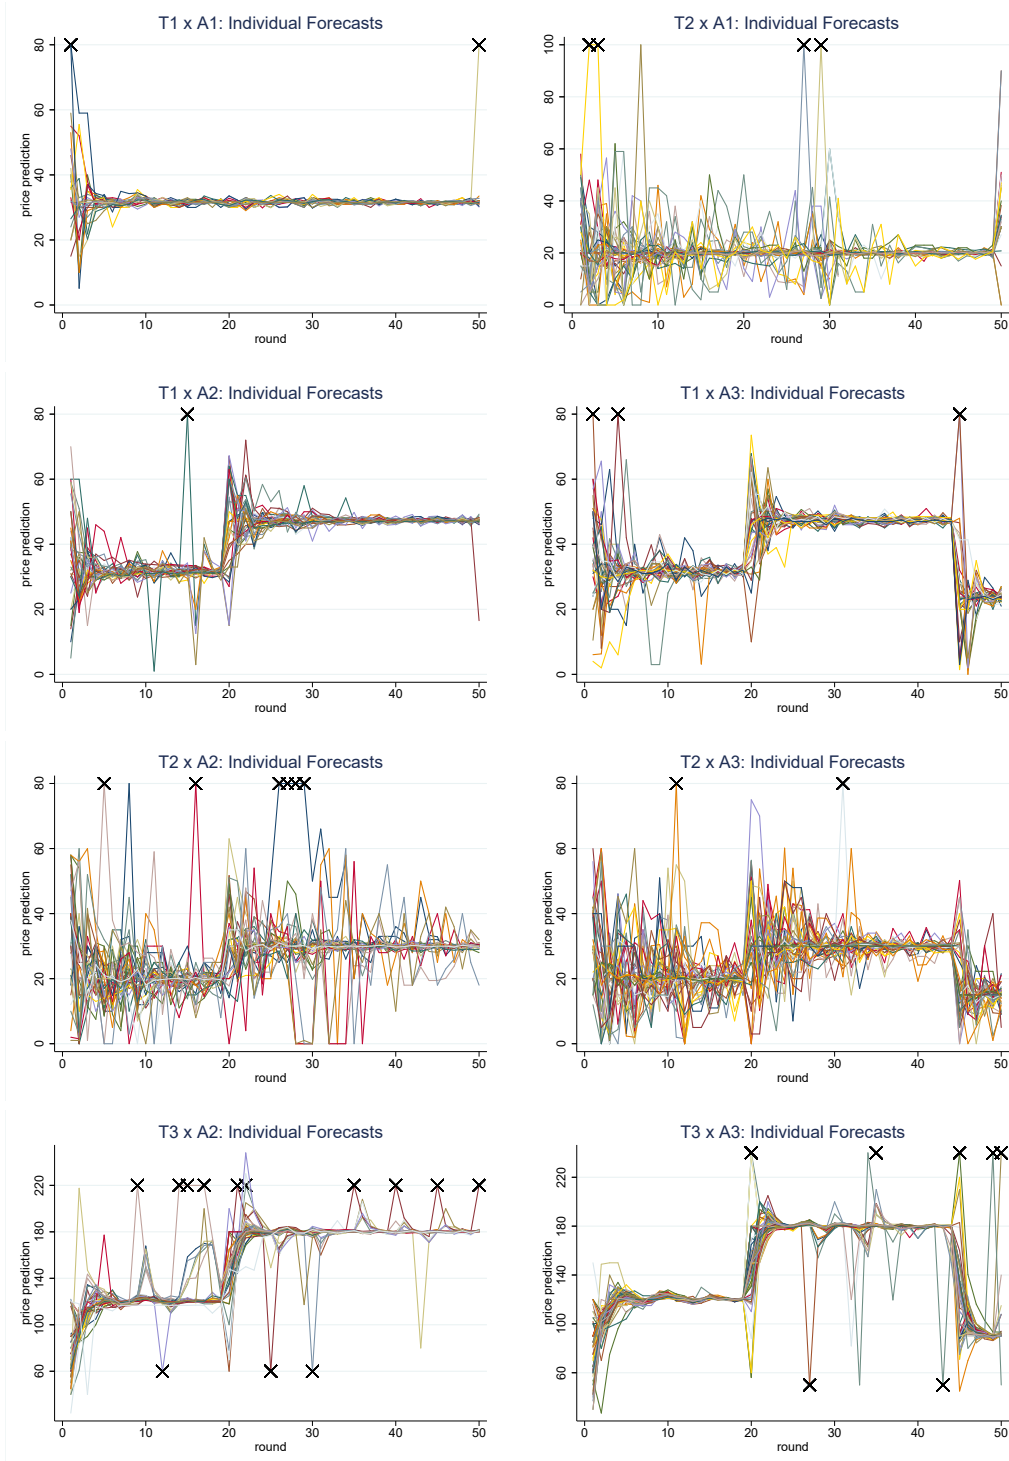


Figure A5: Individual participant predictions



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

Table A2 provides an overall breakdown of the data, including the outliers, to provide a sense of how far away most forecasts are from the model predictions. The table shows three measures of the root squared difference between a subjects submitted forecast and the nearest level-k model implied forecast, where the level-k forecasts are constructed using the standard assumptions given in Section 5.2 in the main text. The root mean squared error/difference (RMSE) for the classifications are quite large. This is almost entirely due to outliers and a minority group of the submitted forecasts. The root median squared error/difference (RMedSE) shows that the majority of forecasts are with one unit of a level-k forecast overall and within 4 units in announcement rounds. The final statistic reported in the table is the 70th percentile of root squared differences. This statistic is chosen because we found that approximately 70% of participants chose a level-k forecast in an announcement round when we use a cutoff value of ± 4.5 for pooled data (see Table 2 in the main text). The column illustrates a treatment-by-treatment breakdown of that classification.

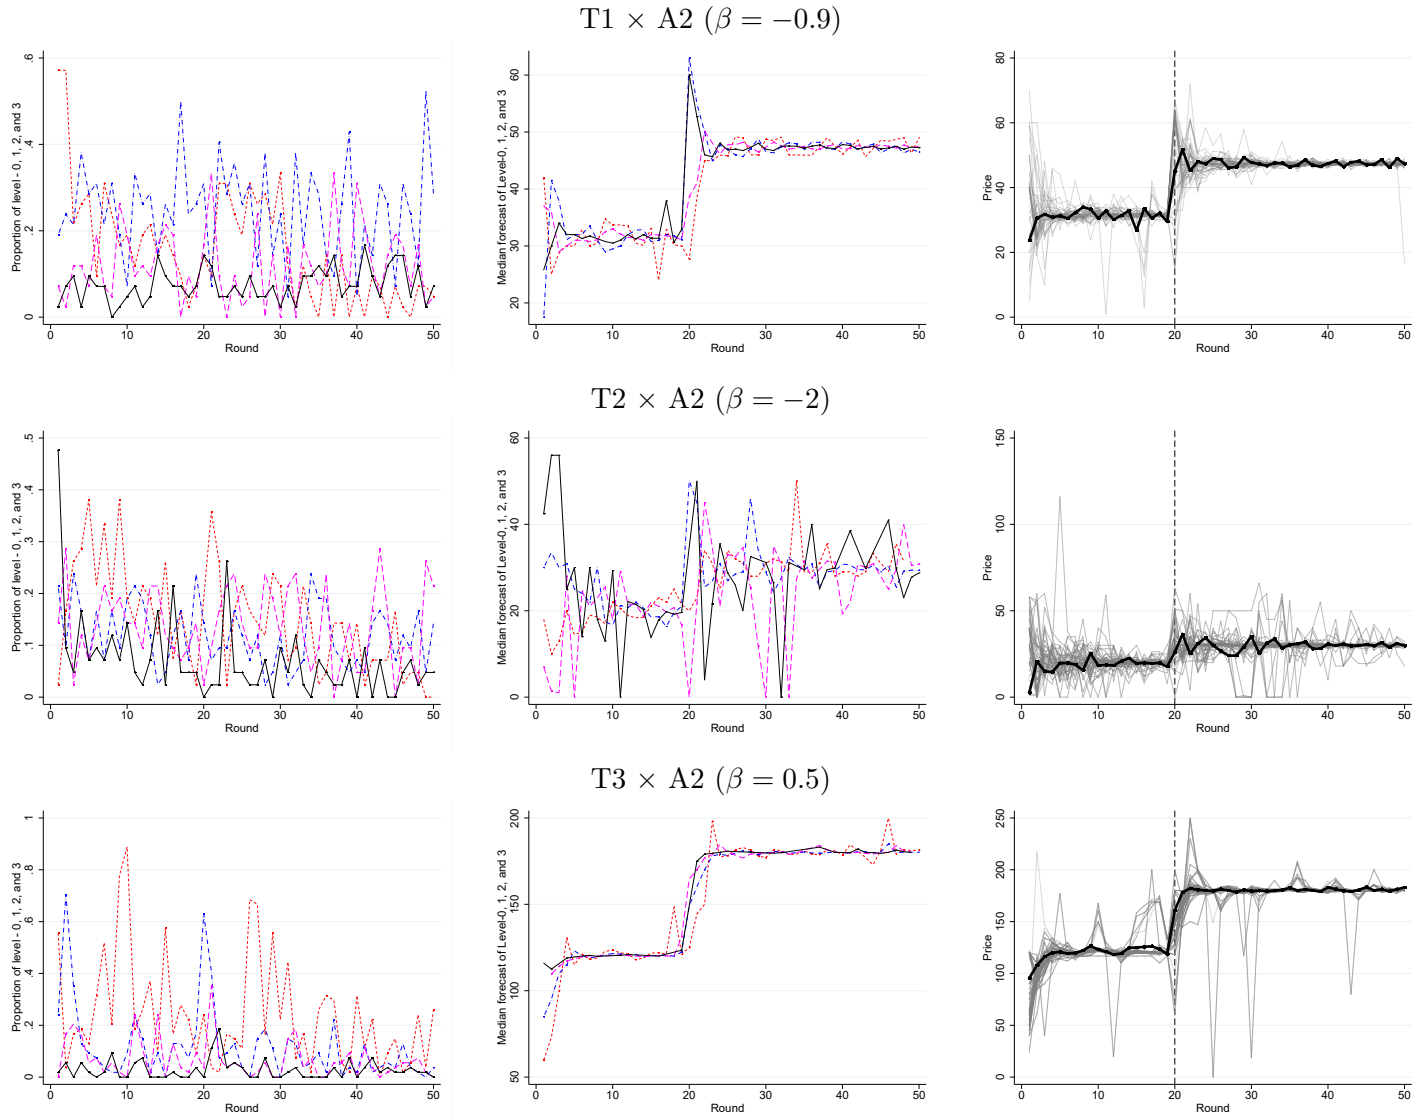
Table A2: Classification of predictions using counterfactual forecast rules

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

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

move the results in favor of the unified model.

Figure A6: Comparing the unified model to experimental data



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

Figure A6 provides the same data breakdown for A2 treatments that we provided for A3 treatments in Figure 2 in the main text. We find similar results here. We identify heterogeneous forecasts that display level-k depths of reasoning in announcement rounds with median individual and mean market dynamics closely matching what was predicted in Section 4.3 in the main text.

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

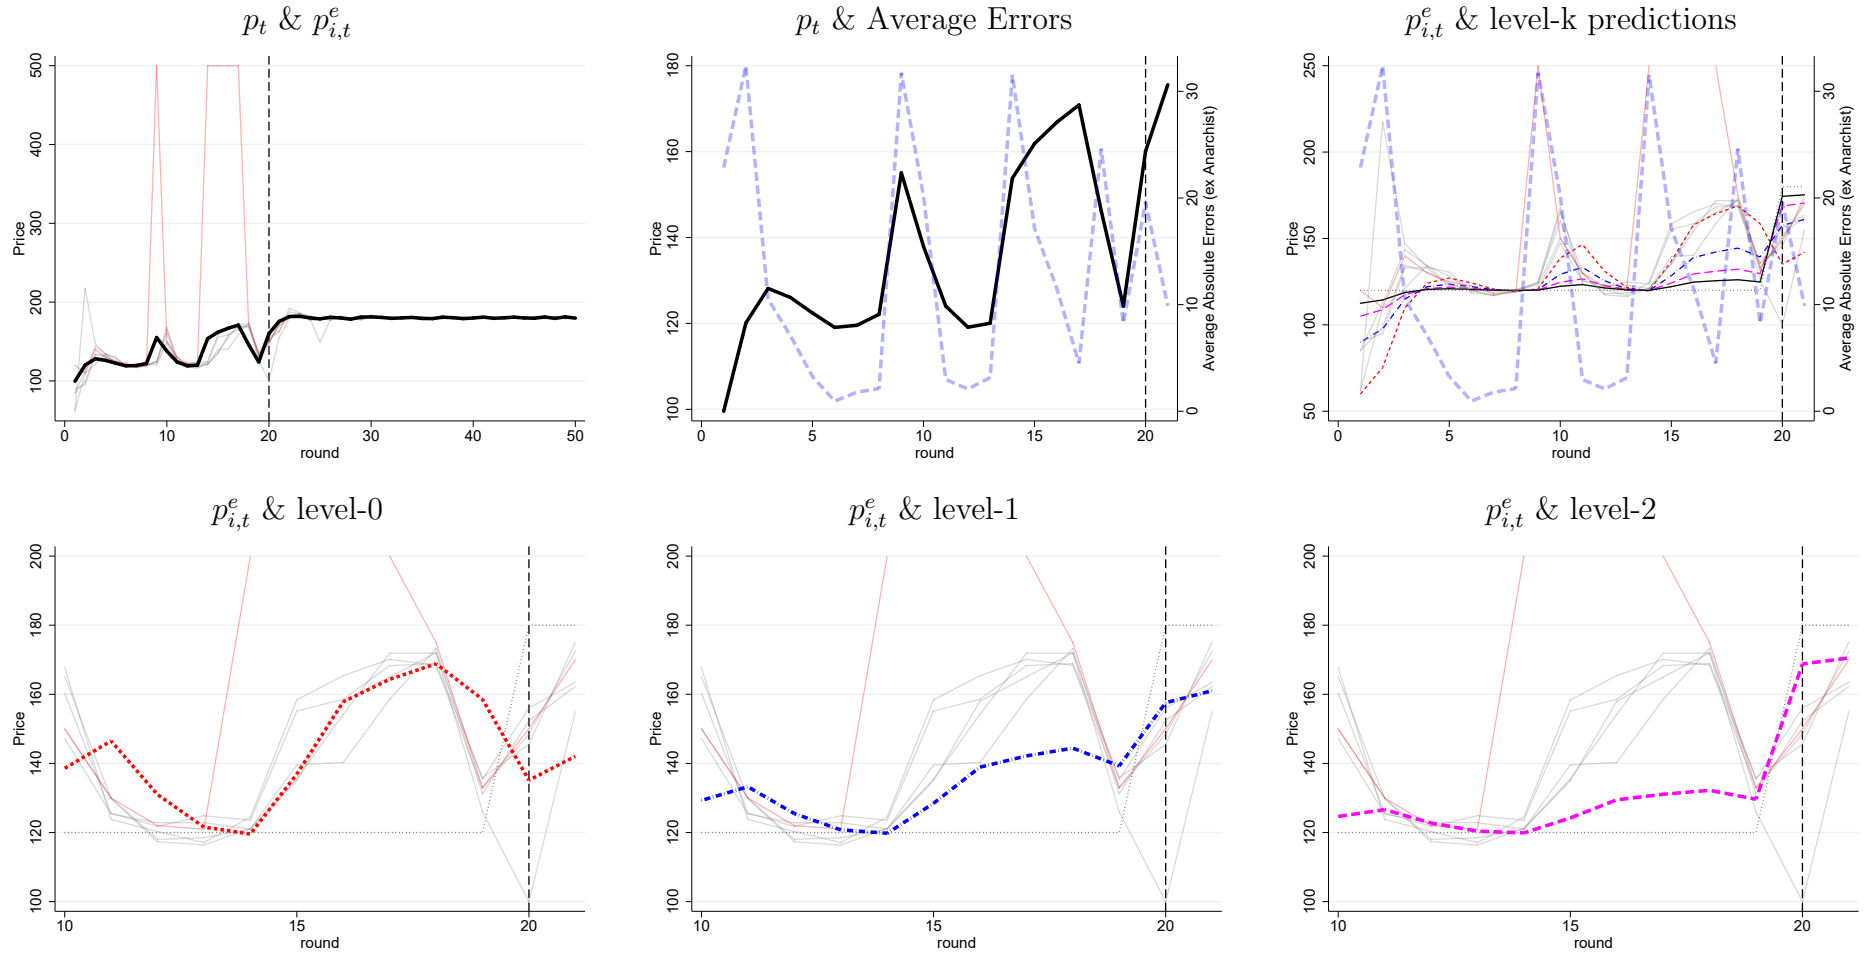
Second, in the T3 treatment we observed one market that had a significant departures in price from the REE after a period of convergence to the REE. You can see the individual forecasts in the third graph on the right of the last row. There is a group of individual forecasts that rise for many periods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

A4 AN ANARCHIST ANECDOTE

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

The final figure in the top row of A7 layers onto the individual expectations the implied level-0, 1, 2, and 3 forecasts using our standard assumptions from Section 5.2. The bottom row of figures in A7 zooms in on the period of interest and plots the implied path of a single level-k forecast on each graph for clarity. It is immediately apparent that each large forecast error generates a shift in behavior by the non-anarchist players. Each shift in behavior is well-captured by one of the level-k deductions.

Figure A7: An Anarchist Anecdote



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

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

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

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

Figure A8 zooms in even further on just rounds 20 and 21 and classifies the individual forecasts types using the method described in Section 5.2 in the main text. Between the two rounds of play, those subjects whose forecasts were closest to the actual price, i.e. experienced the smallest errors, stick with the level-1 forecast. Those subjects who experience larger errors clearly revise up their depth of reasoning, where a revision to level-2 corresponds to what would have been the best forecast to play in round 20 given what occurred. This behavior is consistent with the assumptions that underlie the replicator dynamic’s reflective process that we assume for the unified model.

A5 SUPPLEMENTARY RESULTS FOR SECTION 5


This section provides more detail on the experiment and the results.

A5.1 EXPERIMENT DESCRIPTION

The experiment used a computer based market programmed in oTree.⁷ The market setup follows Bao and Duffy (2016) with additions that accommodate our novel elements. Laboratory participants were randomly assigned to groups of six subjects to form markets. Laboratory participants were told that they are acting

⁷See Chen, Schonger and Wickens (2016) for documentation.

Figure A8: An Anarchist Anecdote Announcement Round



Notes: Individual classified price forecasts in a T3×A2 treatment. The classifications are made by comparing the forecasts to different implied level-k forecasts. The closest implied forecast type determines the classification (see Section 5.2).

as expert advisers to firms that produce widgets. Participants were led through a tutorial that describes the market environment including the exact demand and supply equations that govern the price. Participants were informed that the price depends on the average expected price of all advisers in the market and that prices are subject to small white noise shocks.

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

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

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

we get that price (p) is

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

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

Participants were not able to continue with the experiment until the question was answered correctly.⁸ A worked version of this problem with different numbers was also provided on the printed instruction sheet. The question was designed to verify that each participant knew how to use the equations without teaching the person to solve for the REE. The tutorial and printed instructions are available in the Appendix A10.

Figure A9 shows the graphic user interface (GUI) that participants interacted with during the experiment. The market information is shown in the top right

⁸Four out of 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.

Figure A9: Screenshot of experimental market GUI


Market Game

What is your prediction?

Submit

Period: 10 of 50

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

 Total Earnings: **\$1.94**


corner of the screen. A time series plot of the price and the participant's predictions is provided on the bottom right. A table with the past prices, predictions, forecast errors, and the forecast's earnings is provided on the left-hand-side of the screen.

The payoff function for the participant's predictions is

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

where p_t is the actual market price in the round, $E_{t-1}p_t$ is their prediction for the price in round t , and 0.50 and 0.03 are measured in cents. Negative quantities receive zero cents. The function is presented and explained to participants as part of the tutorial and is the same across treatments. Forecasts must be within 4 units of the actual price to earn money for a forecast. We chose this specification to give participants a high incentive to be precise in their predictions when confronted with announcements. Previous studies have employed point systems that compensate more generously for poor forecasts. For example, in Bao and Duffy (2016) participants needed to be within 7 units to earn points, which ranged from zero to 1300.

In addition to performance pay, subjects received a \$5 show-up fee. In the T2 treatments, subjects also received an additional \$5 of guaranteed compensation to offset the lower earnings that we expected (and which did occur) in these treatments due the difficulty in coordinating.⁹ The difference in guaranteed pay and the treatment settings were not disclosed to the subjects in advance.

⁹Ethics requirements placed on the study mandated that participant payments were on average \$15 AUD per hour.

Announcements for the changes in γ were introduced using a pop-up box. The pop-up box described the change in parameters and participants were required to close the box before they could continue. The announcement would also appear, highlighted in red, across the top of the screen in the announcement period. The information in the top right corner of the GUI would also reflect the change. A minimum price of 0 and a maximum price of 500 was enforced as well. The top bound was not advertised to participants but if chosen a pop-up box would appear informing them of the bound.

A5.2 ADDITIONAL CONVERGENCE RESULTS FOR 5.2.1

To quantify the speed of convergence, we make use of the experimental design where announcements destabilize the market and set off a new period of convergence. This roughly doubles our sample to 111 distinct market periods to study. We measure convergence using three different metrics. First, because of the random noise component of price, we define a round to be converged when the price is within plus or minus three of the steady state price. Based on this cutoff, we simply count the number of rounds in a given interval in which price satisfies this criterion. Columns 2 - 4 of Table A3 shows the count data for the three feedback treatments, where we look at various intervals over the first 19 rounds for all treatments and the comparable intervals for rounds 21 through 38 for treatments with an announcement in period 20. We say that a collection of consecutive rounds has converged if at least 85% of the rounds satisfy the above convergence criterion. Bolded values in Columns 2 - 4 of the table indicate failure to converge. By this metric, none of the feedback treatments (T1, T2, and T3) show convergence within the first five periods of the experiment or within five periods after the first announcement. Convergence is achieved though for T1 and T3 treatments over rounds 6 to 10, rounds 26 to 30, and overall for the full intervals. For the T2 treatments, the 85% threshold is never reached.

The second metric we use to assess convergence is the mean difference in the market price from steady state over the same intervals used for the first metric. Columns 5 - 7 of Table A3 show the mean difference and the t-statistics for a test of the null hypothesis that the mean difference is less-than-or-equal to 3. Bolded values indicate a one-sided rejection of the null hypothesis with a p-value smaller than 0.15. By this metric, convergence is achieved in the T1 treatments within 5 rounds of an announcements and maintained through all other intervals. Convergence is achieved for the T3 treatments in rounds 6-10, but within five rounds after the first announcement. A t-test of the difference in this measure for rounds 2 through 19 versus 21 through 38 confirms that market prices are closer to steady state after the first announcement (bottom row of Table A3) than at start of the experiment, which indicates faster convergence after the announcement. The T2 treatments again show a different pattern. With this metric we only find marginal convergence for rounds 11 - 19 and 31 - 38 in treatments with an announcement. But we do find that prices are on average closer to steady state following the announcement.

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

Table A3: Convergence of price to REE in experimental markets

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

Bolded values do not meet our criteria for market convergence.

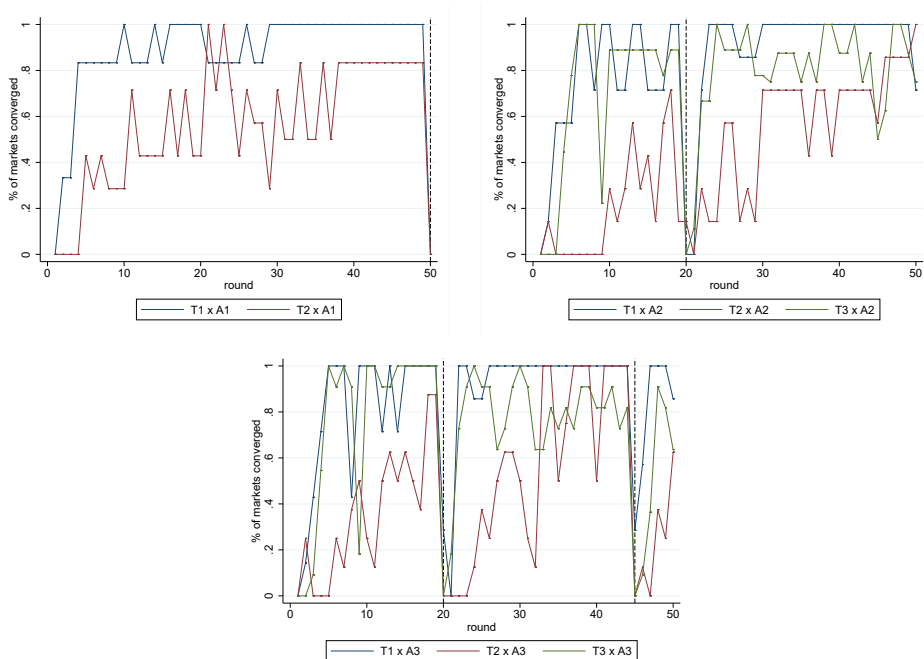
Notes: The table reports three measures of market convergence. Columns 2-4 report the number of rounds where we observe the market price is within ± 3 of the REE price. Columns 5-7 report the mean difference between the market price in a round relative to the REE price for the indicated interval of rounds. Columns 8-10 report the mean earning by participants per round over the indicated interval. The maximum earnings in a round is \$0.50.

all market participants are making accurate forecasts. The last three columns of Table A3 show the mean earnings and the t-statistics for a test of the null hypothesis that average earnings are greater-than-or-equal-to \$0.40. This is our strictest measure of convergence. By this measure, we only observe convergence in rounds 11 - 19 and 31 - 38 for T1 and T3 treatments. We never observe convergence for the T2 treatments.

Figure A10 provides further visualization of the difference in convergence across feedback treatments. The plots show the proportion of markets that recorded two consecutive periods within the plus or minus three band for a rolling window. The T1×A1 and T2×A1 treatments are particularly informative on convergence here. These treatments were designed to closely replicate the main treatment of Bao and Duffy (2016). We replicate their main results by finding relatively quick and maintained market convergence for the T1 treatment and slow or no convergence for the T2 treatments. The T2 treatment also provided the greatest heterogeneity for the speed of convergence with some markets never converging – a phenomena not observed for any market in the T1 and T3 treatments – and some converging quite quickly.¹⁰ The heterogeneous outcomes for the T2 treatments are predicted by the unified model. Small changes in initial conditions to any number of different parameters of the model can lead to coor-

¹⁰One case of non-convergence is particularly noteworthy. One of the T2 x A1 experiments collapsed completely at period 36. One participant was frustrated with the lack of earnings and requested to leave the experiment. All participants in this market were paid their show-up fee and earnings up to period 36 and then dismissed from the experiment.

Figure A10: Convergence of price to REE in experimental markets



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

dination on the REE or to a market that completely destabilizes. Further, the unified model correctly predicts that all T1 and T3 treatments should robustly converge.

A5.3 ADDITIONAL LEVEL-K RESULTS FOR 5.2.2


Figure A11 shows histograms of individual forecasts in round one and in each announcement round for each feedback treatment. The gray bars show the model implied level-k forecasts with ± 3 band. The T2 round 20/50 predictions provides the clearest level-k deductions because the large negative feedback ($\beta = -2$) makes each level-k prediction very distinct. While there is not strong evidence for level-k reasoning in the first period,¹¹ this changes once participants have played multiple rounds and an announcement occurs. For these announcement rounds, Figure A11 and Table 2 (main text) show a majority of participants playing level-k or the high level-k/REE forecasts.

For the robustness exercises regarding the definition of the level-0 forecast see Section A6

For the results on oscillating deductions see Section A7

¹¹This lack of evidence may reflect difficulty in establishing level-0 forecasts in round one. Also, some participants do not appear to understand the game's structure in the first period.

Figure A11: Laboratory subjects' forecasts in announcement rounds



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

A5.4 ADDITIONAL REVISION AND LOSS RESULTS FOR 5.2.3

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

To test the three features of the replicator dynamic, we make use of the announcements in the A2 and A3 treatments. The announcement rounds provide a clear intervention from which to identify level-k deductions. They generate large forecast errors for many participants, and they provide distinct counterfactual level-k predictions, which we can use to identify subsequent revisions to the depth of reasoning in the experimental data. Specifically, comparing individual outcomes and predictions in the announcement rounds to the round following the announcement, we can assess who has revised their depth of reasoning, how the revision compares to the best level-k forecast one could have chosen in the announcement period, and whether those who changed strategy experienced larger forecast errors. To maximize the data and to not exclude those who decided to switch from a non-classified strategy to a level-k strategy, we do not impose a cutoff when classifying a person's forecast as level-0, 1, 2, 3, or the REE for this analysis. Classifications are made based on whichever level-k strategy the submitted forecast is closest to in mean squared error.

Table A4 reports the results for the first and second announcements across all treatments. The first column shows the proportion of individuals who, conditional on changing strategies, are classified as selecting the best counterfactual strategy from the previous period, which was often a lower level of reasoning than the one played in the announcement round as predicted by the unified model. The second column reports the proportion of participants whom we identify as not changing their strategy. The remaining columns report the difference in mean absolute forecast errors experienced by changers and non-changers and the deliberation time when selecting their new forecast.

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

Table A4: Revisions and loss

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

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

A5.5 ADDITIONAL LEVEL-K DYNAMICS RESULTS FOR SECTION 5.2.4

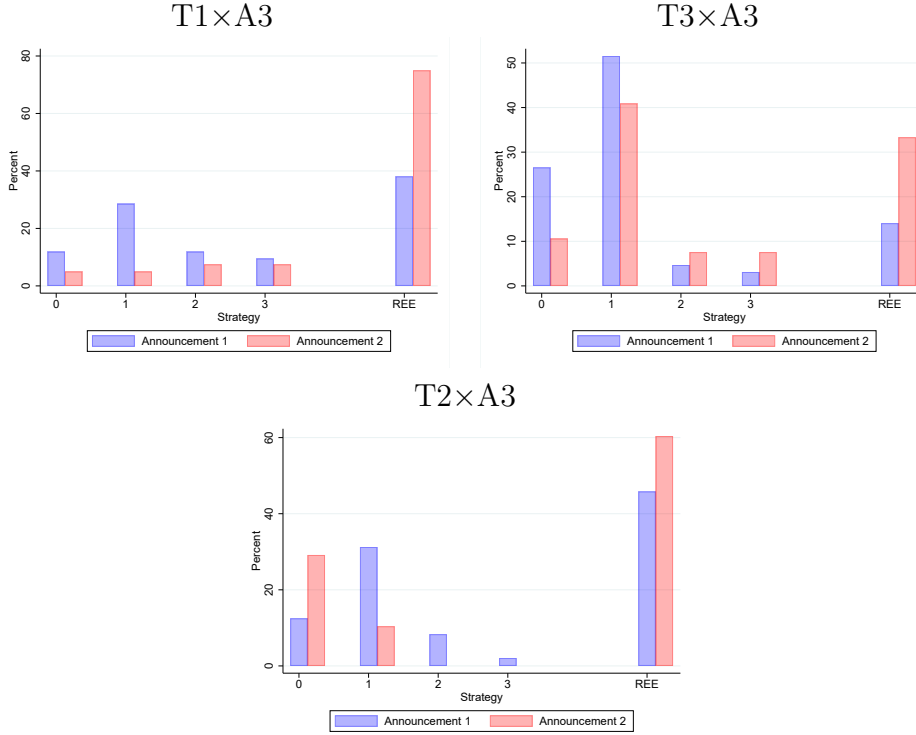
The unified model also predicts that when $|\beta| < 1$ we should see increasing depth of reasoning over time during periods when the market structure is constant. We can test this prediction by looking at the distribution of strategies that are played across the same subjects in the A3 treatments with two announcements. The unified model predicts that over time more people will select higher level-k forecasts for the T1 and T3 treatments, but not for the T2 treatments. Figure A12 shows the distribution of forecasts for levels-0 to 3 and REE. We can see for the T1 and T3 cases that the distribution shifts to the right. More subjects choose higher-level forecasts, or are consistent with the REE forecast in the second announcement than in the first. We find that a Kolmogorov-Smirnov equality of distributions test rejects the null of equality at the 5% level for the T1 and T3 treatments. The T2 treatments, however, shows a different result. For T2 treatments, we observe a bifurcation in which subjects either choose a low levels of reasoning or they jump to the REE.

A5.6 ADDITIONAL QUANTITATIVE RESULTS FOR SECTION 5.2.5

We fit the model to the experimental data at the market level. Table 3 in the main text averages over the individual market outcomes from the same treatments. Table A5 shows the underlying data from each market.

Each model that features heterogeneous types is initialized to the first realized price and to the distribution of level-k types observed in period one for each market. Afterwards, the model makes predictions based solely on the evolution of price, adaptive learning, or the replicator, depending on which model is used.

Figure A12: Increasing or decreasing depths of reasoning over time



Notes: The distribution of classified forecast types observed in the experimental treatments with two announcements.

The learning model starts initial beliefs at the average of the individual forecasts in period one. After period one it updates according to the evolution of data implied by the model and beliefs for the chosen gain. The simulated data is compared to experimental data and the mean squared error is calculated.

Each model is optimized individually by searching over a grid of gains $\phi \in [0, 1]$, or replicator parameters $\alpha \in [0, 2]$, or both in the case of the unified model. The optimal coefficients are shown in Table A6. Both the replicator and adaptive learning are required to best fit the data in T1 and T2 treatments. In many of the T3 treatments, however, naive expectations and fixed level-k reasoning is chosen as the best model. This reflects the fact that many markets converge very quickly to steady state, but not as quickly as RE implies. This is also reflected in the results for the adaptive learning case where a naive model is found to best fit the data for all markets. In subsequent exploration, which is not shown here, we have found that a $\phi > 1$ plus level-k reasoning is preferred. That is consistent with a trend following behavior similar to what many other positive feedback experiments have found.

A5.7 ADDITIONAL DISCUSSION

The experimental evidence provides strong support for Hypothesis 1 (stability). Large negative feedback results in slow convergence, or nonconvergence, to the REE price, while convergence is achieved for $|\beta| < 1$. In addition, the speed of convergence measured in multiple ways appears to increase following an announcement treatment (see Table A3). Increases in convergence speed in treat-

Table A5: MSE between experimental data and competing models

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

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

Table A6: Parameter estimates of competing models

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

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

ments T1 & T3 are also supported by the increase in the depths of reasoning we observe among subjects when there are multiple announcements: see Figure 2.

We find strong support for Hypothesis 2 (level-k reasoning). We observe level-k deductions taking place in each of the announcement treatments with clear bunching around the k-level predictions in the histograms shown in Figure A11. Comparing the individual forecasts to the model implied forecasts in announcement rounds, we classify between 50% and 70% of subjects, depending on the chosen cutoff, as Level-0, 1, 2, 3, or REE. Our classifications also coincide well with the deliberation times we observe among participants, with level-0 participants spending less time deliberating than level-3.

We find support for Hypothesis 3 (replicator dynamics). Focusing again on announcement periods, we find that some fraction of subjects are classified as using the same depth of reasoning in the announcement period and in the period following the announcement. These subjects on average had lower forecast errors in the announcement period than those subjects who appear to change strategies, and they spent less time deliberating in the next round. In addition, for those we classify as changing their strategy, we find evidence that a high proportion are changing to the best strategy (see Table A4). As predicted by our theory, many of those changes correspond to decreases in the subject’s k-level depth of reasoning.

Finally, we find mixed evidence for Hypothesis 4 (level-k dynamics). We observe revisions over time in depth of reasoning for the T1 and T3 treatments. There were also more high level-k forecasts played for second announcements compared to first announcements, along with quicker convergence (see Figure 2 and Table A3). In addition, we do observe a bifurcation in the distribution of classified strategies played in the T2×A3 treatments between the two announcement rounds with more level-0 and REE forecasts played in the second announcement round. The reduction in the depth of reasoning in favor of level-0 forecasts observed here is consistent with hypothesis 4. However, the increase in the fraction of people who choose the REE forecast is at odds with the unified model. This finding also explains why the RE forecast fits the aggregate price data fairly well in the quantitative evaluation of competing models discussed in the the previous sub-section.

We speculate that the high proportion of REE forecasts observed in the T2×A3 treatment’s second announcement round may be due to the fact that in the experiment negative prices are not allowed. In an announcement round, many high-level forecasts predict either 0 or γ in the T2 treatment. Therefore, a subject’s menu of forecasts has a finite number of distinct choices. With finite choices and bounded prices, its plausible that some subjects will engage in sufficient reflection to engender more coordination on the REE, which is in the interior of the price space. This is an interesting avenue for future research.

A6 ROBUSTNESS: LEVEL-0 FORECAST DEFINITION

To classify the types of forecasting strategies that participants use, we must assume a shared level-0 forecast. Our baseline assumption is that level-0 is a two-round moving average of past prices. To demonstrate that our results are robust to this assumption, we conduct two exercises. First, we replicate the results in Table 2 and A4 of the main text using a four round moving average as the shared

level-0 forecast. Second, we study how overall classification of types and of the level-0 type changes when we assume last periods price as the level-0 forecast, a two-period moving average, a four-period moving average, or three different constant gain specifications.

Table A9 replicates Table 2 with the four period average. The number of people we classify as level-k reasoners increases slightly under this definition overall. The regression estimates are mostly unchanged. We retain statistical significance for the hypothesis test conducted on announcement rounds with a comparable F-stat obtained to the original specification.

Table A7: Classifying participant's forecasts as Level-k - Robustness Check

Within ± 3 of Level-k in announcement rounds				Differences in deliberation time (seconds)		
	1	20/50	45	Variable	[1]	[2]
Total Classified	47.3% [33.9% , 57.0%]	65.8% [50.6% , 73.6%]	66.0% [49.4% , 69.9%]	Level-0	-9.73 (1.113)	-1.97 (0.672)
Level-0	14.8% [11.0% , 15.1%]	7.8% [4.31% , 9.48%]	5.1% [4.49% , 6.41%]	Level-1	-8.61 (1.144)	-0.16 (0.629)
Level-1	7.3% [6.45% , 8.60%]	25.0% [19.3% , 27.6%]	14.1% [14.1% , 14.1%]	Level-2	-5.50 (1.312)	-0.24 (0.991)
Level-2	6.5% [1.89% , 6.45%]	5.2% [4.31% , 5.75%]	3.8% [1.92% , 3.85%]	Level-3	-6.67 (1.332)	-0.33 (1.054)
Level-3	3.2% [1.07% , 1.13%]	3.2% [2.58% , 3.74%]	4.5% [3.21% , 5.13%]	Level-0 x Ann	47.52 (8.623)	3.58 (6.065)
REE	15.6% [13.4% , 15.6%]	24.7% [20.1% , 27.0%]	38.5% [25.6% , 40.4%]	Level-1 x Ann	45.47 (4.881)	11.90 (4.741)
				Level-2 x Ann	8.61 (9.058)	11.08 (8.546)
				Level-3 x Ann	63.82 (11.92)	23.02 (8.260)
N	372	348	156	Cons	41.16 (0.526)	112.52 (4.227)
Hypothesis tests of deliberation time regressions				Individual FE	yes	yes
				Round FE	no	yes
$H_0 : \text{Level-0} - \text{Level-3} = 0$				R-squared	0.030	0.253
$H_0 : (\text{Level-0} \times \text{Ann}) - (\text{Level-3} \times \text{Ann}) = 0$				N	18,367	18,367

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

Table A8: Revisions and loss - Robustness Check

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

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

Table A8 replicates Table A4 for the four-round average level-0 assumption. The results are slightly stronger on all categories relative the previous definition.

Table A9 shows the classification results for the ± 3 cut off for different level-0 assumptions. In general, the proportion of subjects that we classify as level-k forecasters of any type increases as we consider level-0 forecasts with longer averages or weighted averages of past observed prices. Therefore, the results we provide as our baseline are the most conservative estimates we obtained.

Table A9: Classifying participant’s forecasts as Level-k - Robustness Check

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

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

A7 OSCILLATING DEDUCTIONS WITH STRATEGIC SUBSTITUTES

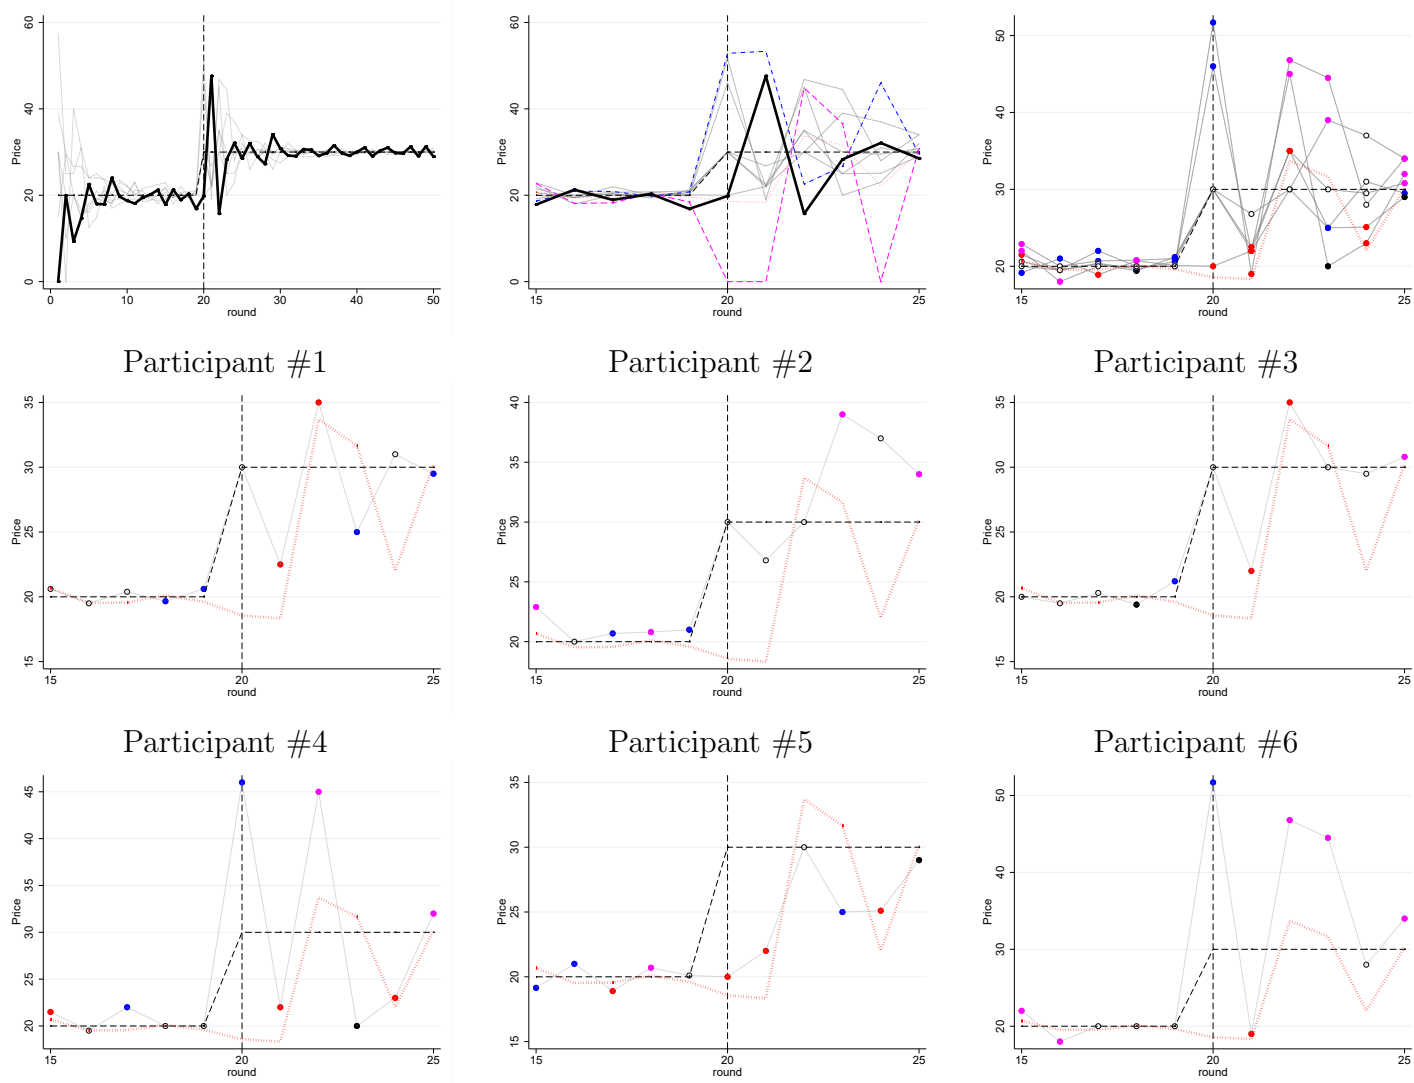
Our experimental results also shed light on an element of level-k reasoning that recently has been called into question in the literature. García-Schmidt and Woodford (2019) and Angeletos and Sastry (2021) put forward models of bounded rationality that modify level-k reasoning to rule out oscillating deductions when there is negative expectational feedback. Angeletos and Sastry (2021) writes,

“We are not aware of any experimental evidence of this oscillatory pattern. We suspect that it is an unintended “bug” of a solution concept.”

In our experiments, we provide evidence that oscillating deductions occur. Figure A13 is based on a T2×A2 treatment. The NW panel shows market price in bold and players’ forecasts in light gray; the N panel shows market price in bold and the model implied level k forecasts (note truncated time-span around announcement). The NE panel shows players’ forecasts with dots colored to identify level-k classification. Remaining panels provide model implied level-0 forecasts in dashed red and player forecasts colored to identify level-k classification.

From Figure A13 is evident that some individual participants make oscillating deductions over time. Obvious examples include participants #1 and #4. Indeed, four out of six of the participants clearly make oscillating deductions with forecasts above and below the REE after the announcement occurs in round 20. The oscillations occur despite the experience of the price not oscillating for many periods prior to the announcement. This experience of tranquility combined with how close many of the forecasts are to level-k deductions suggests that participants contemplated oscillations consistent with classic level-k reasoning. Moreover, they took action with money at stake which was consistent with such deductions.

Figure A13: Oscillating Deductions: Individual forecasts from experimental market with treatment T2 ($\beta = -2$)



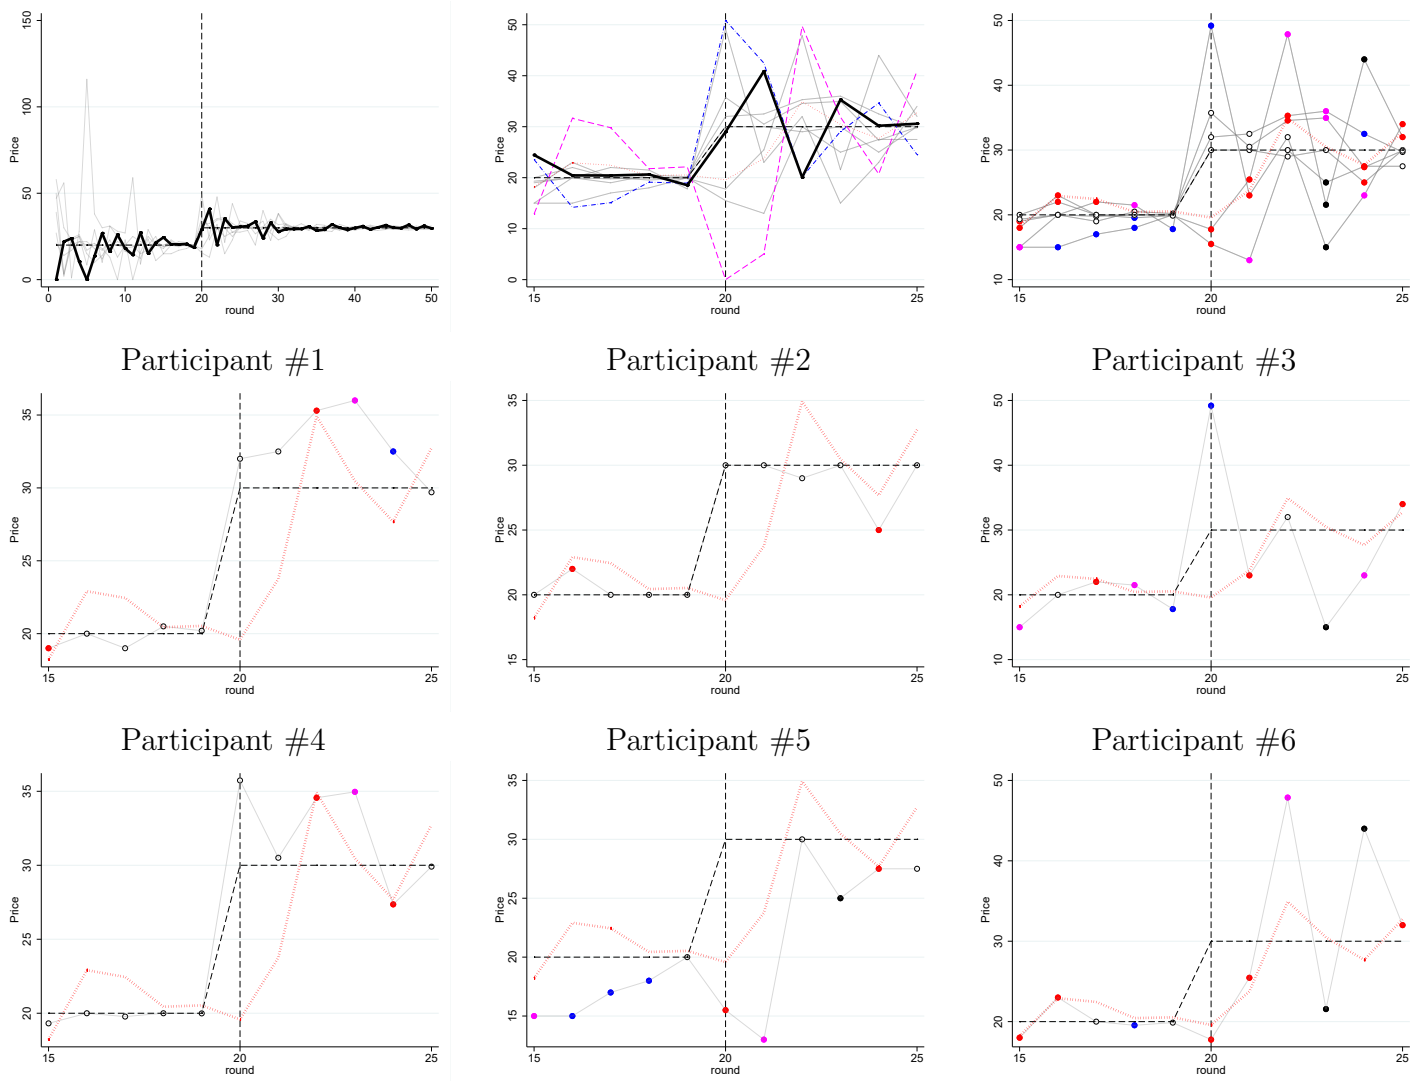
Notes: The first plot shows all individual forecasts and the market price from a single market. The second plot shows the model implied level-k forecasts and the data in a window around the announcement. The remaining figures classify each of the forecasts as a level-k type, which is indicated by the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

Figures A14 shows an additional example of a market with clear oscillating behavior and Figures A15 shows one market with more monotonic behavior.¹²

To further illustrate point, Figures A16 show the same type of analysis applied to a T3 treatment where $\beta = 0.5$. Level-k deductions do not imply oscillations in this case and indeed none are observed. Individual forecasts conform nicely to level-k deductions based on our proposed level-0 forecast. This suggests that people do not abandon level-k deductions in environments with strategy substitutability. Level-k deductions describes forecasting behavior in our experiment when strategic actions are both compliments and substitutes.

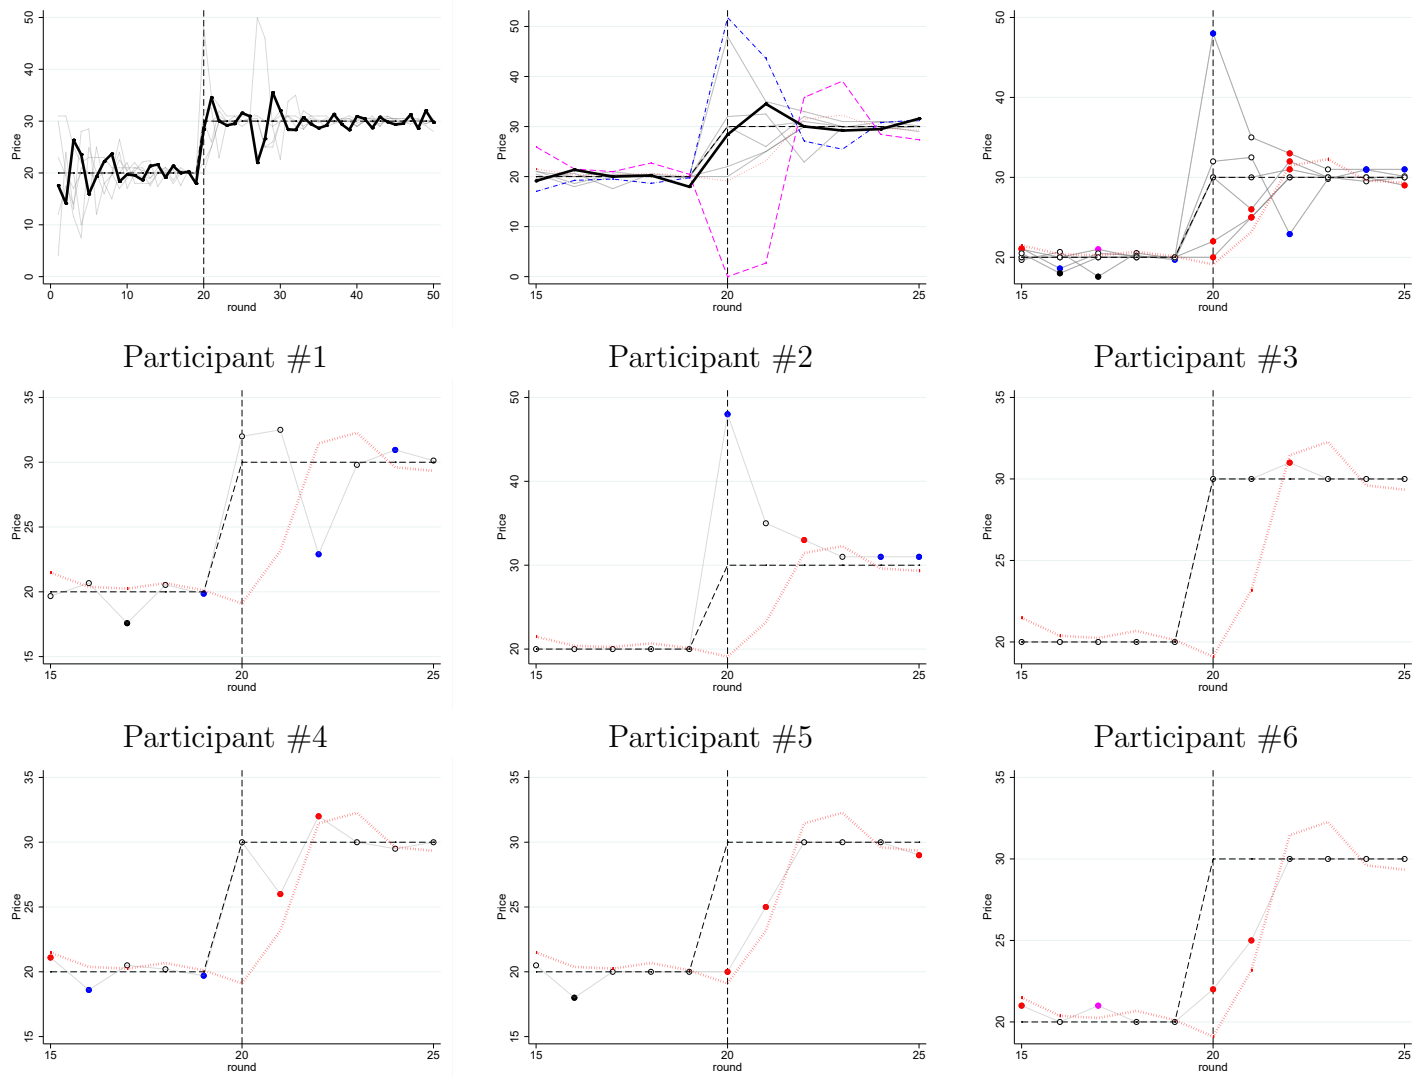
¹²We do not plot the level-3 forecast in the middle figure because it makes the graph harder to interpret by requiring a larger scale of the y-axis. For the markets we show, no one chooses it in the announcement round. This of course is not true in general. We observe people choosing exactly level-3 deductions in some markets as can be seen in Figure 2 in the main text.

Figure A14: Example 2: Individual forecasts from experimental market with treatment T2 ($\beta = -2$)

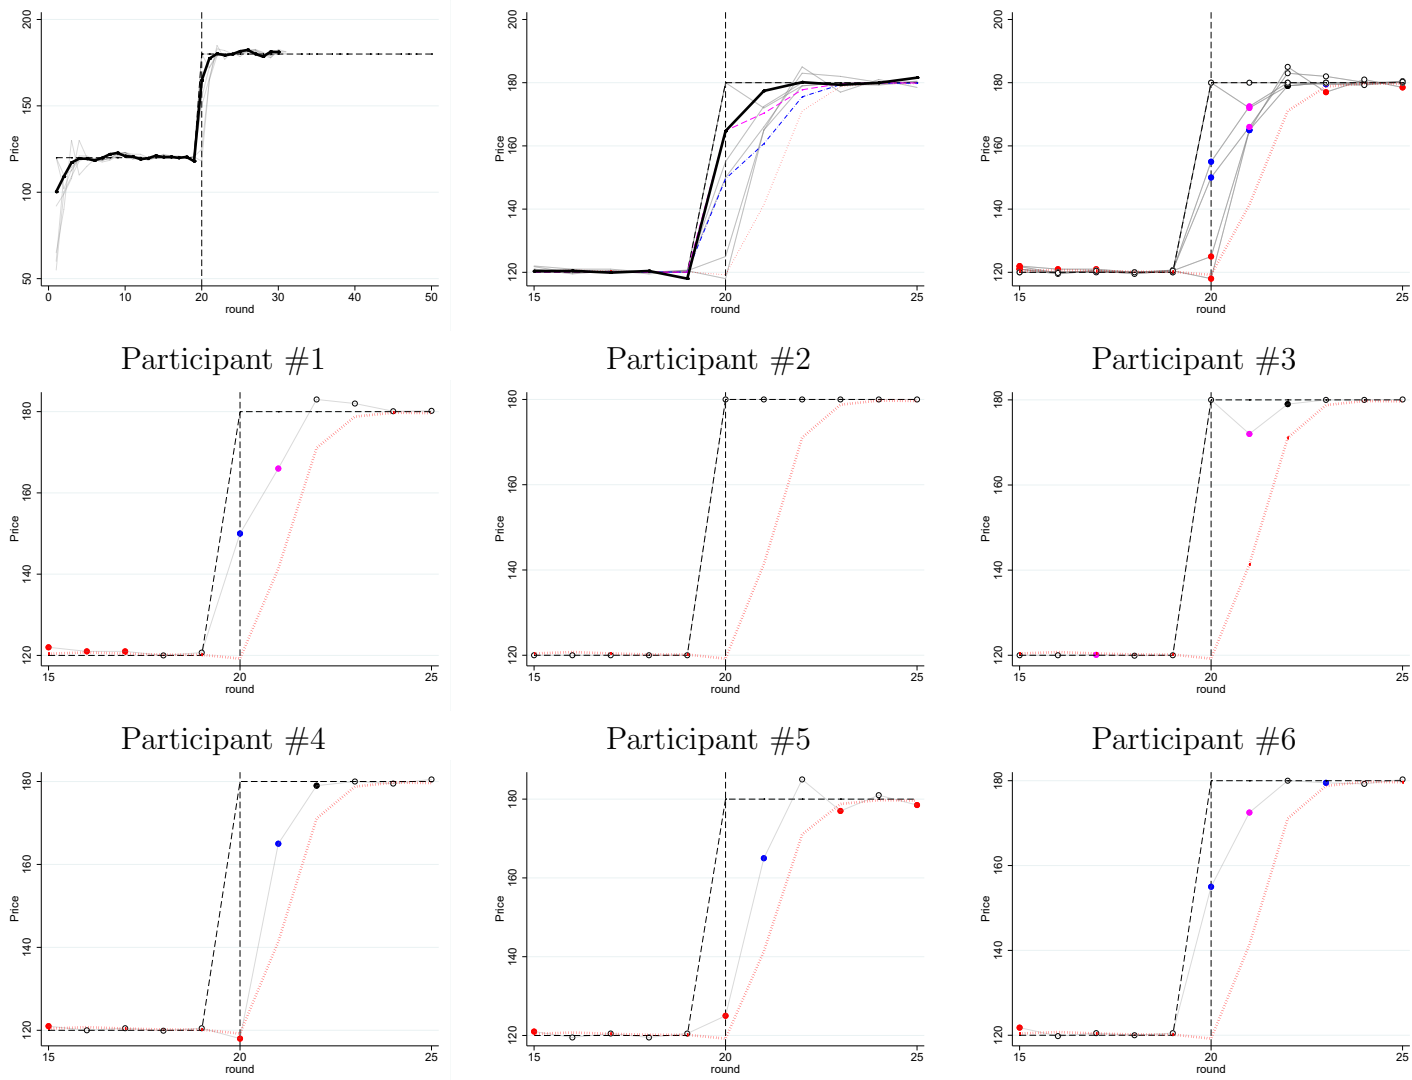


Notes: The first plots shows all individual forecasts and the market price from a single market. The second plot shows the model implied level-k forecasts and the data in a window around the announcement. The remaining figures classify each of the forecasts as a level-k type, which is indicated by the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

Figure A15: Example 3: Individual forecasts from experimental market with treatment T2 ($\beta = -2$)



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

Figure A16: Example 4: Individual forecasts from experimental market with treatment T3 ($\beta = 0.5$)

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

A8 EXIT SURVEY RESULTS

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

A8.1 EXIT SURVEY QUESTIONS

1. Please rank the importance of each option below to the formation of your price forecast in each period:
 - a. The history of market prices
 - b. The market equations
 - c. The history of my own price forecasts
 - d. The history of my own forecasts errors
 - e. My expectation about the average price forecast in the period

2. Please rank the importance of each option below to the formation of your price forecast **following the announcements**:
 - a. The history of market prices
 - b. The market equations
 - c. The history of my own price forecasts
 - d. The history of my own forecasts errors
 - e. My expectation about the average price forecast in the period

3. Which of the following statements best describes your thinking before making each forecast?
 - a. I looked at the past prices and made my best guess based on their recent movements. I never used the equations.
 - b. I made a guess about what the average forecast might be based on past prices and then used the equations to determine my own forecast using that guess.
 - c. I made a guess about what the average forecast might be and used the equation to work out the price only when I did a poor job of forecasting in the previous round. Otherwise, I just looked at past prices and made my best guess.
 - d. I made a guess about what the average forecast might be and used the equation to work out the price only when there was an announced change in the market. Otherwise, I just looked at past prices and made my best guess.

4. Please rank the importance of each option below to **other participants**, which you believe they may have used to make their price forecasts:
 - a. The history of market prices
 - b. The market equations
 - c. The history of their own price forecasts
 - d. The history of their own forecasts errors
 - e. Their expectation about the average price forecast in the period
5. Please rank the importance of each option below to **other participants**, which you believe they may have used to make their price forecasts **following the announcements**:
 - a. The history of market prices
 - b. The market equations
 - c. The history of their own price forecasts
 - d. The history of their own forecasts errors
 - e. Their expectation about the average price forecast in the period
6. If you do not feel like the strategy you used was well-captured by the survey questions, then please use this box to explain your strategy

A8.2 EXIT SURVEY RESULTS

Survey questions (1), (2), (4), and (5) used a drop-down menu with options: “very important”, “somewhat important”, and “did not consider.” Table A10 and Table A11 shows the cumulative importance of each factor where “very important” is assigned a zero, “somewhat important” is assigned a one, and “did not consider” a two. Therefore, the lower the value, the more important the information. Consistent with level-k reasoning, we find that on average subjects rated the equations and the forecast of the average expectation as more important to their own forecast than they believed it was to others. This is consistent with a belief that others are less sophisticated. We observe the results on the full sample and when restricting to only people who played a level-k forecast in the announcement periods with the ± 3 cutoff. The latter consistently rank the equations as important to them than they are to their perceived competitors, which is consistent with the level-k assumption that others players are perceived as less sophisticated.

Table A10: Tabulated survey results for Q1 and Q4

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

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

Figure A17 shows the responses to question 3 separated by treatment. The most common response is (b), which is:

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

This response is consistent with level-1 behavior.

Table A11: Tabulated survey results for Q2 and Q5

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

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

A9 UNIFIED DYNAMICS IN THE NK MODEL

Consider the following RE forward model:

$$y_t = \gamma_t + \beta E_t y_{t+1}, \quad (\text{A14})$$

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

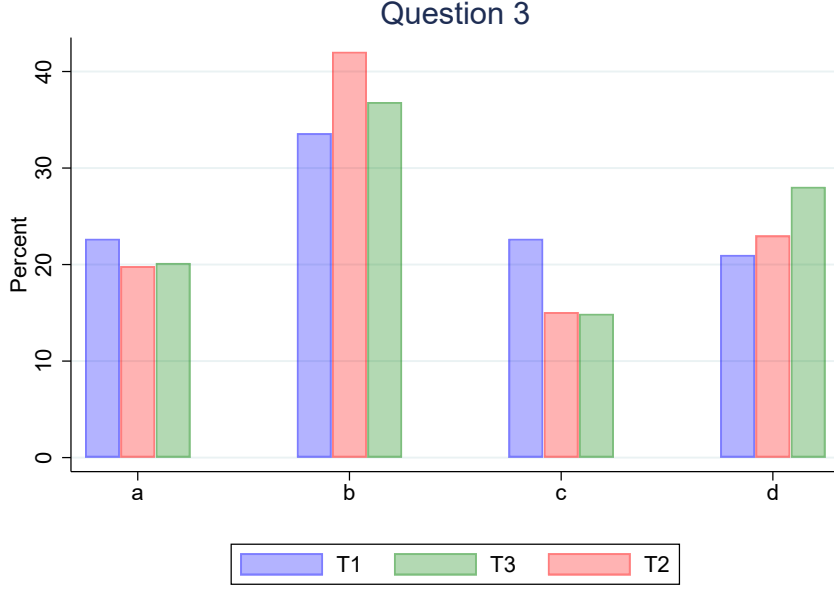
$$y_t = \gamma_t + \beta \sum_k \omega_{kt} \cdot E_t^k y_{t+1}, \quad (\text{A15})$$

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

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

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

Figure A17: Exit survey question 3 responses



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

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

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

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

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

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

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

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

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

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

Application to Bilbiie's NK model

Recall our laboratory NK model:

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

$$\pi_t = \xi E_t \pi_{t+1} + \kappa y_t \quad (\text{A21})$$

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

To map this model into the general theory developed above, we proceed as follows: writing $y_t = (x_t, \pi_t)$ (of course, vectors are always columns), let $z_t \in \{S, F, N\}$ be the underlying state, so that $r_t = r(z_t)$ and $i_t = i(z_t)$. Then $z_0 = S$ and the transition for z_t is

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

Set

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

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

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

Thus

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

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

$$a_{zt} = a_{zt-1} + \chi_z(z_t) \cdot \phi \cdot (y_{t-1} - a_{zt-1}),$$

and

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

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

The REE of the model is obtained by backward induction. Let y_S and y_F be the output-gap/inflation pairs associated with γ_S and γ_F respectively. Since

$y = 0$ when $\gamma = \gamma_N$ we can solve for y_S and y_F using backward induction. We have

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

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

A10 EXPERIMENT MATERIALS

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

NEGATIVE FEEDBACK CASE

Computer based tutorial:

- What is your role?

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

- What makes you an expert in this market?

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

- What is a widget?

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

- Why do the firms need to forecast the price?

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

- How am I paid?

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

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

where p is the actual market price, and 0.50 and 0.03 are measured in cents. If your forecast is off by more than 4, you will receive \$0.00 for your forecast. Therefore, you will receive \$0.50 for a perfect forecast, where $p = \text{your price forecast}$, 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

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

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 $p = \text{your price forecast}$, 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

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

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:

REFERENCES

- Angeletos, George-Marios and Karthik A Sastry**, “Managing Expectations: Instruments Versus Targets,” *The Quarterly Journal of Economics*, 2021, 136 (4), 2467–2532.
- Bao, Te and John Duffy**, “Adaptive versus eductive learning: Theory and evidence,” *European Economic Review*, 2016, 83, 64–89.
- 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.
- 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.
- 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,” *American Economic Review*, 1992, 82 (5), 1254–1278.

Widget Game Instruction Summary:

- Your job is to forecast the price of a widget next period
- Demand for widgets is determined by the market price
 - $q = A - B p$
- The total supply of widgets to the market is determined by the **average** of all price forecasts submitted to the market
 - $q = D \times \text{average price forecast}$
- Combining supply and demand, we have the **key formula** that determines price in the market
 - $P = A/B - D/B \times \text{average expected price} + \text{noise}$
 - Recall that *noise* is small and **on average equal to zero**
- **An Example:** $A = 120$, $B = 2$, $D = 1$, and $\text{noise} = 0$, what is price if the average price forecast is 42?
 - $p = 60 - \frac{1}{2} \times \text{average price forecast}$
 - $P = 60 - \frac{1}{2} \times 42 = 60 - 21 = 39$
- You are paid based on **accuracy of your forecast** according to the following formula
 - $\text{Payment} = 0.50 - 0.03 (p - \text{your price forecast})^2$
 - 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
 - $q = C + D \times \text{average price forecast}$
- Combining supply and demand, we have the **key formula** that determines price in the market
 - $P = (C - A)/B + D/B \times \text{average expected price} + \text{noise}$
 - Recall that *noise* is small and **on average equal to zero**
- **An Example:** $A = 0, B = 2, C = 60, D = 1$, and $\text{noise} = 0$, what is price if the average price forecast is 42?
 - $p = 30 + \frac{1}{2} \times \text{average price forecast}$
 - $P = 30 + \frac{1}{2} \times 42 = 30 + 21 = 51$
- You are paid based on **accuracy of your forecast** according to the following formula
 - $\text{Payment} = 0.50 - 0.03 (p - \text{your price forecast})^2$
 - 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