

Adaptive Heuristics in Macroeconomics*

Cars Hommes^{a,b}

^aBank of Canada

^bUniversity of Amsterdam

July 31, 2025

Abstract

Abstract. This paper discusses behavioral and experimental macroeconomic research on expectations and learning with simple adaptive heuristics. Agents may use heuristics when there is limited information and/or when the environment is too complex to fully understand. Emphasis will be given to how expectations feedback affects the choice of heuristics and the (in)stability of the adaptive heuristics learning process of heterogeneous agents. Finally, we discuss how policy can manage expectations and shift adaptive heuristics of heterogeneous agents from destabilizing trend-following strategies to adaptive expectations to stabilize a complex adaptive heuristics macro system.

JEL Classification: E7, E32, C92, D84, D83

Keywords: heuristics, expectations, adaptive learning, behavioral and experimental macroeconomics.

*This paper has been written for an invited session on *Macroeconomics and Information* at the Econometric Society World Congress, 18-22 August, 2025, Seoul, South-Korea. I am grateful to the session organizer Yuri Gorodnichenko for an opportunity to discuss this work at the ESWC. Opinions expressed in this paper are those of the author and do not necessarily reflect those of the Bank of Canada or its staff. Any remaining errors are mine.

1 Introduction

This paper discusses recent work on adaptive heuristics in macroeconomics. A *heuristic* is a simple decision rule, while *adaptive heuristics* refers to a learning process of (the parameters of) heuristics or to the selection process to choose from a pool of competing, heterogeneous heuristics. Our focus here will be on *forecasting heuristics* as a way of modeling expectations as an adaptive learning process of heterogeneous agents in complex macroeconomic systems.

Our paper is not the first to stress the importance of simple heuristics in models of bounded rationality in macroeconomics. For example, Simon's (1955) work on bounded rationality advocated the use of a satisficing heuristic, a decision-making strategy that aims for a good enough solution rather than the optimal one. Sargent (1993) is an early stimulating study of bounded rationality models in macro, where agents behave as econometricians or as artificial agents. More recently, Moll (2024) argues for non-rational expectations in heterogeneous agents macroeconomic models and points to simple heuristics as one of the promising ways to go. Heuristics have a long history in psychology, and their importance for economic decision making has been stressed by Tversky and Kahnemann (1974), who argued that [op cit Moll (2024)]:

“people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations”. [hereafter added from Tversky and Kahnemann (1974)] *“In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors.”*

The added quote relates directly to the Lucas critique that simple forecasting heuristics lead to biases and systematic forecasting errors and that agents will learn from these mistakes and adapt their behavior accordingly. But this argument perhaps has been overemphasized. There are simple heuristics that perform well in some environments.

For example, the naive expectations heuristic, using the last observation as the forecast, is optimal when the forecast variable follows a random walk. Therefore, an AR(1) near unit root forecasting heuristic is nearly optimal in a near unit root environment, and is in fact hard to beat by more complex forecasting rules. Nevertheless, a resolution of the Lucas critique requires a more realistic model of an adaptive learning process among different heuristics. We will see that some forecasting heuristics are consistent with empirical realizations and may in fact be hard to beat by more complex rules. This is similar to Gigerenzer and Todd (1999), who stressed the importance of “*simple heuristics that make us smart*”, where agents learn to use efficient, fast and simple procedures that fit the environment.

Broadly, two different types of learning of heuristics have been proposed in the literature. The first is adaptive learning of simple forecasting heuristics, such as AR(1) or AR(2) rules. These are simple, parsimonious forecasting models that are generally misspecified and give rise to non-rational but near-rational, learning equilibria (Bullard, 1994; Evans and Honkapohja (2001), Branch, 2006, Hommes and Zhu, 2014). The second is a heuristics switching model (HSM), where agents may switch between heterogeneous heuristics based upon their relative performance (Brock and Hommes, 1997; Anufriev and Hommes (2012)). Heuristics have been tested extensively in laboratory experiments with human subjects; see Hommes (2021) for a survey. Whether a forecasting heuristic works well may depend on the environment. In particular, we will argue that the type of expectations feedback is crucial for the success of different heuristics.

Heuristics are useful when information is limited, or when the environment is too complex to fully understand or to compute more complex (rational) forecasts. As we will see, even with limited information simple heuristics may sometimes lead to the full information rational equilibrium (FIRE). In such case, coordination on such a simple heuristic may explain why the system converges to FIRE. In general, heuristics are important as a coordination device, because coordination of a large group of agents on a simple heuristic

seems a more plausible outcome of a learning process than coordination on a very complex rule (see e.g. Arifovic et al., 2019 for a simple lab experiment, where subjects learn simple heuristics (a steady state or a 2-cycle) in a complex environment.

Our work relates to recent work on survey data analysis by Weber et al. 2025, showing how a changing inflation environment alters the learning process of individuals. They conclude that the strength of the response of expectations to exogenously provided information speaks directly to the inattentiveness of individuals to such news. We argue how a changing feedback alters the learning process of individuals and the aggregate outcome (Hennequin et al. 2025, see also Section 9). Different feedback alters individual learning heuristics and as a consequence the aggregate macro behaviour.

Heuristics are also related to limited information equilibria. Limited information equilibria are game theory concepts that describe how players behave when they have limited information about each other or the game. These equilibria can occur in games with incomplete information, also known as asymmetric information. Young (2008) describes adaptive heuristics in game theory as simple behavioral rules that are directed towards payoff improvements, but may be less than fully rational. Examples include reinforcement learning (Erev and Roth, 1998) and regret matching (Hart, 2005).

The paper is organized as follows. We start with heuristics in the cobweb framework (Section 2), one of the models that Muth (1961) used to introduce rational expectations. Section 3 focuses on inflation dynamics in a New Keynesian Philips Curve model where agents learn an optimal simple AR(1) forecasting rule. Section 4 discusses heuristics in asset price experiments and Section 5 presents the heuristic switching model (HSM) and fits it to experimental data. The role of positive versus negative expectations feedback is discussed in Section 6, while Section 7 presents genetic algorithm learning of optimal anchor and adjustment heuristic in positive versus negative feedback environments. Section 8 uses the New Keynesian framework to show how the Central Bank can manage expectations and trend-following behavior through more aggressive inflation targeting. Section 9

discusses some recent results of how (in)stability of expectations in the lab depends on eigenvalues of the expectations feedback system and how policy can shift expectations to the stable region. Section 10 summarizes some conclusions about adaptive heuristics in macroeconomic expectations.

2 Cobweb model

The stability of equilibria in an economy depends on the way agents form expectations. Lucas (1986) pointed out that stability is not just a matter of a theory of adaptive learning, but should be viewed as an experimentally testable hypothesis: “*But to be useful, stability theory must be more than simply a fancy way of saying that one does not want to think about certain equilibria. I prefer to view it as an experimentally testable hypothesis, as a special instance of the adaptive laws that we believe govern all human behavior.*” (Lucas, 1986, pp. S424 - 425).

In response to Lucas’ stability hypothesis a large literature on learning-to-forecast (LtF) experiments has emerged, pioneered by Marimon, Spear and Sunder (1993) and Marimon and Sunder (1993); Adam (2007) is one of the first LtF experiment in a macro DSGE setting; see Hommes (2021) for a review of experimental macroeconomics. In a LtF experiment human subjects have to forecast a price p_t that depends on the forecasts of a group of subjects. For example, in a simple LtF experiment the realized market price may depend on the average forecast \bar{p}_t^e of a group of subjects, that is,

$$p_t = f(\bar{p}_t^e). \tag{1}$$

Learning to Forecasts Experiments (LtFEs) may be viewed as repeated Keynesian beauty contest games introduced in Nagel (1995); see the survey of Mauersberger and Nagel (2018). In a LtF experiment subjects’ *only* task is to forecast prices, with all other behavior in the economy computerized by (rational) theory through the mapping f . Often, the

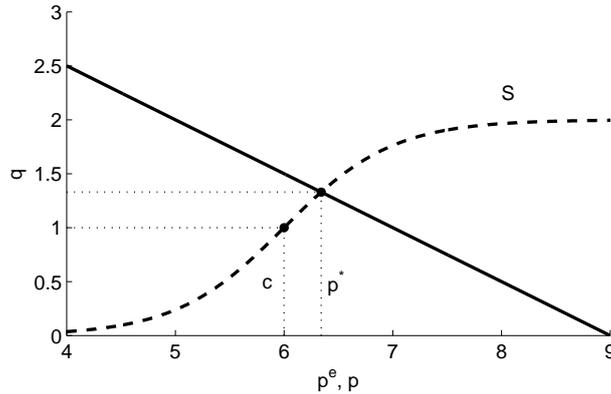


Figure 1: Cobweb model with linear demand and nonlinear, but increasing, underlying the LtFE in HSTV (2007)

payoffs in an LtFE are proportional to the forecasting accuracy. In most LtFEs subjects do not know the exact underlying equilibrium equations, but only have *limited, qualitative information* about the market or the economy. For example, subjects know that price is determined by market clearing and they know about the sign, but not the magnitude, of the expectational feedback. Typically, subjects have no information about the behavior of other subjects and only observe past prices and their own past predictions and earnings.

In this paper we will compare LtFEs in different frameworks: the cobweb model (this Section), the asset pricing model (Section 4) and the New Keynesian macro model (Section 8). As we will see, there are two important cases to distinguish; *positive feedback*: the mapping f in (1) is increasing; and *negative feedback*: the map f is decreasing. Negative feedback arises in the cobweb model of a perishable consumption good, where high expectations lead to high production and low market prices. In contrast, positive feedback dominates in speculative markets, where optimistic expectations fuel asset demand leading to higher prices. Macro models also often exhibit positive feedback, for example, from inflation expectations to realized inflation.

Figure 1 illustrates a cobweb model with linear decreasing demand and a nonlinear, S-shaped, increasing supply curve, as used in the LtFEs of Hommes et al. (2007) (HSTV2007

hereafter). For the cobweb model the mapping f in (1) is derived from market clearing, with J producers, and is given by

$$p_t = D^{-1}\left(\sum_{j=1}^J S(p_{j,t}^e)\right) = \frac{a + \epsilon_t - \sum_{j=1}^J S(p_{j,t}^e)}{b}. \quad (2)$$

Here, producer j has price expectations $p_{j,t}^e$, the demand curve is assumed to be linear with slope $-b$ and subject to small shocks ϵ_t and supply is nonlinearly increasing.

The rational expectations solution is a fixed point $p^* = f(p^*)$ of the map f or equivalently an intersection point of demand and supply curves. In the LtFE, however, demand and supply and their rational intersection point $p^* = 5.93$, are unknown to the subjects.

Important questions then for testing in the lab are: (i) will subjects be able to learn the rational outcome?, and (ii) which heuristics play a role in the adaptive learning process in the lab?

Cobweb model in the Lab

Before discussing what happens in the lab, it is useful to discuss two simple heuristics. Figure 2 illustrates simulations of the realized market prices under two benchmark expectations rules, namely naive expectations ($p_t^e = p_{t-1}$) and sample average ($p_t^e = (1/t) \sum_{i=0}^{t-1} p_i$). In these simulations the cobweb model satisfies $S'(p^*)/D'(p^*) < -1$, so that the steady state p^* is *unstable* under naive expectations. Prices then fluctuate endogenously and converge to a stable 2-cycle (buffered with small noise). Along the 2-cycle forecasting errors of the naive forecast are large and predictable, since they are strongly (negatively) correlated. Therefore, in the cobweb context, the naive heuristic is systematically wrong and therefore subject to the Lucas critique: agents would learn from the 2-cycle and adapt their forecasting behavior. .

The second plot in Figure 2 shows that under sample average learning the price converges to the rational expectations steady state p^* . Stated differently, the RE price is stable under sample average learning.

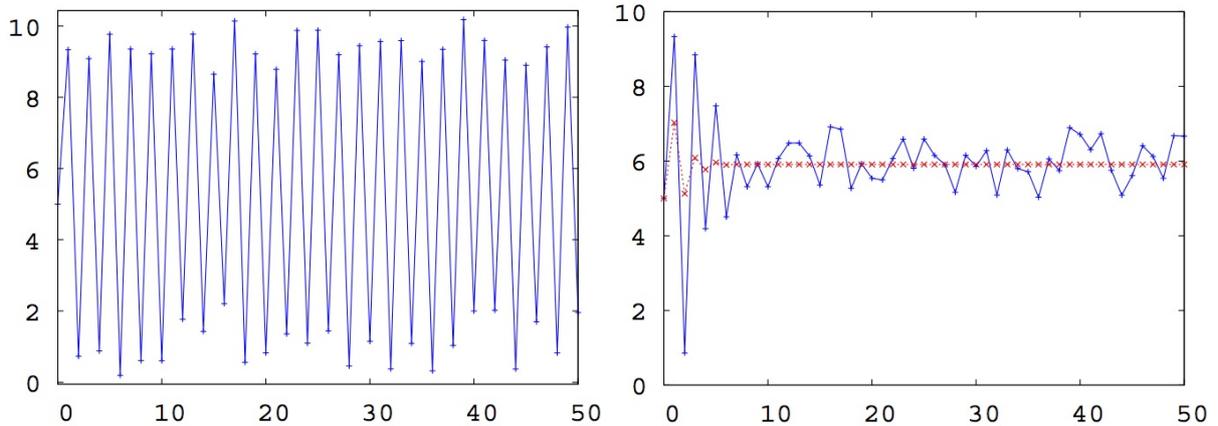


Figure 2: Simulations of an unstable cobweb model under naive expectations (left) and learning by sample average (right plot). Under naive expectations price converges to a stable 2-cycle, with systematic forecasting errors. Under sample average learning price converges to the rational expectations steady state price (with small shocks).

The cobweb model is an example where one simple heuristic, naive expectations, predicts price fluctuations around an unstable steady state, while another simple heuristic, sample average learning, predicts convergence to the steady state. What happens in the lab? Will agents learn from their systematic mistakes and converge to RE?

Figure 3 illustrates two different groups from two different treatments in the cobweb LtFEs in HSTV(2007). In the stable treatment (left plot) parameters of demand and supply are such that $S'(p^*)/D'(p^*) \approx -0.95 > -1$, so that the steady state price p^* is stable under naive expectations. In the stable treatment subjects in the lab are able to learn and prices converge to the RE outcome, with small fluctuations (due to small noise ϵ_t) around the RE steady state. Hence, in the stable treatment in the cobweb model with *limited information*, subjects are able to learn and coordinate on the rational outcome. In contrast, the behavior in the unstable treatment (see right plot in Fig 3) is very different, with irregular large amplitude fluctuations around the RE price. In contrast to the regular noisy 2-cycle fluctuations under naive expectations, (see Figure 2 left plot) in the unstable lab treatment (see Figure 3, right plot) fluctuations are large and irregular and have three characteristics: (i) the mean market price is close to the RE price p^* ; (ii) no

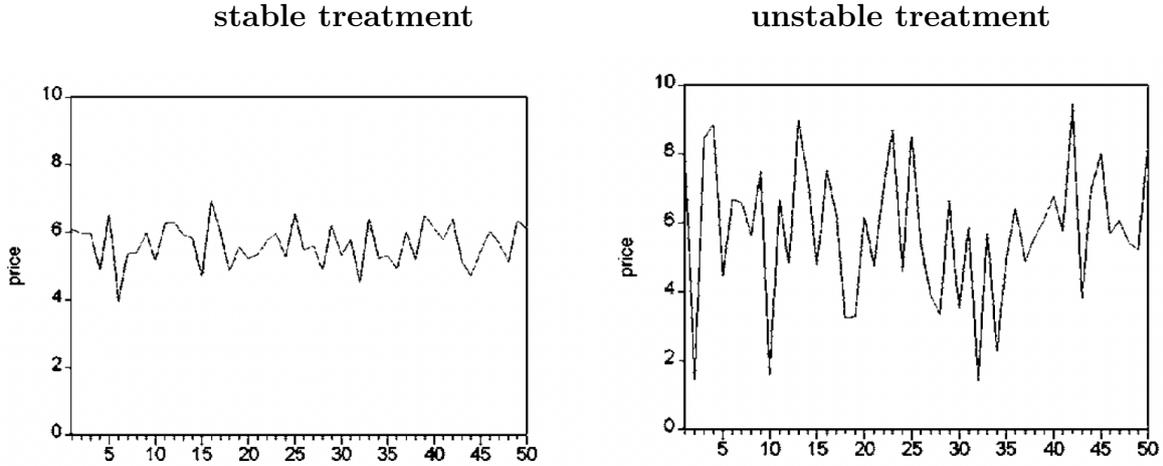


Figure 3: Stable treatment converges to RE. Unstable treatment (right plot): mean p^* (RE) no autocorrelations and excess volatility

autocorrelations in prices, and (iii) excess volatility, that is, much larger fluctuations than under RE.

In order to explain these characteristics, Hommes and Lux (2013) fit a genetic algorithm (GA) learning model to these experimental data, as illustrated in Fig 4. GA agents learn the parameters of a simple heuristic, an AR(1) forecasting rule of the form $\alpha + \beta(p_{t-1} - \alpha)$. In the stable treatment the GA-learning algorithm converges to the RE outcome $\alpha^* = p^*$ and $\beta^* = 0$.

In the unstable treatment the GA-learning algorithm does not converge completely, but fluctuates around the RE outcome $\alpha^* = p^*$ and $\beta^* = 0$. In both the stable and unstable treatments the AR(1) parameter β_t converges close to its RE value 0. The mean parameter α_t fluctuates more around its mean 0 and occasionally deviates substantially in the unstable treatment.

In terms of the Lucas critique, Lucas' prediction was justified, in the unstable cobweb treatment agents do *not* coordinate on a 2-cycle with systematic forecasting errors, but subjects adapt their behavior. The laboratory results are consistent with GA-learning of the parameters of an AR(1) heuristic, with the persistence parameter β_t converging close

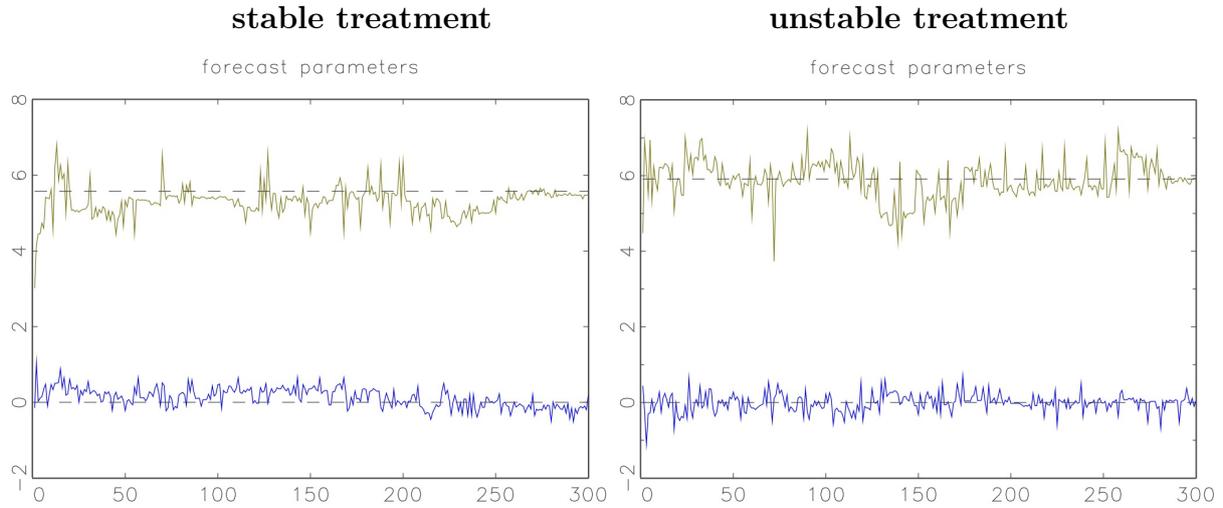


Figure 4: Cobweb simulations under GA learning. Stable treatment (left plot) converges to RE. Unstable treatment: mean p^* (RE), no autocorrelations and excess volatility. Unstable treatment: learning parameters fluctuate around RE values $\alpha^* = p^*$ and $\beta^* = 0$.

to 0, so that (almost) all autocorrelations in prices and forecast errors wash out. The GA-learning model also explains the excess volatility under learning compared to RE, and thus explains the three characteristics observed in the experimental data.

3 Learning an optimal AR(1) heuristic

The naive forecasting heuristic is biased and leads to large and systematic forecasting errors in a cobweb environment, where expectations feedback is negative (i.e., a higher price-forecast yields a lower market price). In macro and finance often expectations feedback is positive— that is, a higher forecast leads to higher prices. This happens, for example, in speculative asset markets, where higher asset price expectations fuels asset demand leading to higher asset prices, or in New Keynesian models of inflation and output, where higher inflation expectations may trigger a surge in inflation.

How does the naive forecasting heuristic perform in positive feedback systems that are common in macro and finance? We will argue that the simple naive heuristic is a fairly good predictor in an economy with (strong) positive feedback. This should not come as a surprise when noting that if prices follow a random walk (a unit root process), the naive forecast is the optimal forecast in terms of mean-squared-error. Hence, in near unit root systems, that are common in macro and finance, the naive forecast is nearly optimal and, therefore, naive expectations is in fact hard to beat by other more complex forecasting models.

Behavioral Learning Equilibrium (BLE)

Hommes and Zhu (2014) introduced the notion of Behavioral Learning Equilibrium (BLE), where agents use an optimal AR(1) rule in a complex and unknown environment. Agents do not know the actual law of motion (ALM) of the economy, but rather believe that prices follow a simple AR(1) process. Hence, agents' PLM is an AR(1) process of the form

$$x_t = \alpha + \beta(x_{t-1} - \alpha) + v_t, \tag{3}$$

where v_t is an IID process. The parameter α is the (long run, unconditional) mean and β is the persistence coefficient (the first-order auto-correlation). Along a BLE, these param-

eters, however, are not free, but must satisfy consistency requirements for two observable statistics: α must equal the (long run) mean, and β must equal the first-order auto-correlation of the (unknown) actual law of motion (ALM) of the economy. The BLE is a simple and parsimonious *misspecification equilibrium* (Branch, 2006), since the PLM is an AR(1) model, which is different from the ALM. Along a BLE, the parameters α^* and β^* are fixed at their optimal values, minimizing the MSE. Therefore, a simple BLE forecasting heuristic may be viewed as a first order approximation and may in fact be hard to beat by more complex forecasting strategies.

But how would agents know or compute the optimal parameters α^* and β^* ? Hommes and Sorger (1998) introduced a simple and intuitive learning process of a BLE: sample auto-correlation learning (SAC-learning), where the parameters are time-varying, α_t equals the sample average and β_t is the first-order auto-correlation, given by¹

$$\alpha_t = \frac{1}{t+1} \sum_{i=0}^t p_i,$$

$$\beta_t = \frac{\sum_{i=0}^{t-1} (p_i - \alpha_t)(p_{i+1} - \alpha_t)}{\sum_{i=0}^t (p_i - \alpha_t)^2}$$

This simple learning mechanism has an intuitive behavioral interpretation, with agents “guestimating” the mean and the persistence coefficient from time series observations.

BLE in the NKPC

Hommes and Zhu (2014) introduced BLE in a very simple setting of a linear univariate model driven by autocorrelated shocks. Applications of this simple structure include the asset pricing model driven by autocorrelated dividends and the New Keynesian Philips Curve (NKPC) with inflation driven by marginal costs. Here we focus on the NKPC (see

¹SAC-learning is similar to OLS learning, as is often used in models of adaptive learning.

Lansing, 2009)

$$\begin{cases} \pi_t = \lambda \pi_{t+1}^e + \gamma y_t + u_t, \\ y_t = a + \rho y_{t-1} + \varepsilon_t, \end{cases} \quad (4)$$

where π_t is inflation, π_{t+1}^e inflation expectations, y_t the exogenous driving variable (output gap or real marginal cost), and u_t and ε_t are independent i.i.d. white noises.

Under rational expectations inflation is a linear function of the driving variable. Therefore, under RE, inflation π_t inherits its properties from the driving variable marginal costs or output gap. In particular, under rational expectations, the persistence (autocorrelation) structure of inflation equals the persistence of its driving variable y_t .

Now consider the NKPC with a linear AR(1) forecast for expected inflation. Given the perceived law of motion (PLM) in (3) the 2-period ahead inflation forecast is

$$\pi_{t+1}^e := \alpha + \beta^2(\pi_{t-1} - \alpha),$$

and the actual law of motion (ALM) becomes

$$\begin{cases} \pi_t = \lambda[\alpha + \beta^2(\pi_{t-1} - \alpha)] + \gamma y_t + u_t \\ y_t = a + \rho y_{t-1} + \varepsilon_t. \end{cases} \quad (5)$$

The consistency requirements for BLE in the NKPC are: (i) the mean α of the PLM equals mean inflation $\bar{\pi}$ of the ALM (5), and (ii) the persistence coefficient (i.e., the first-order autocorrelation) β of the PLM must equal the autocorrelation $F(\beta)$ of realized inflation in the ALM (5).

Sample Autocorrelation Learning (SAC-learning; Hommes and Sorger, 1998) in the NKPC is given by

$$\begin{cases} \pi_t = \lambda[\alpha_{t-1} + \beta_{t-1}^2(\pi_{t-1} - \alpha_{t-1})] + \gamma y_t + u_t \\ y_t = a + \rho y_{t-1} + \varepsilon_t \end{cases} \quad (6)$$

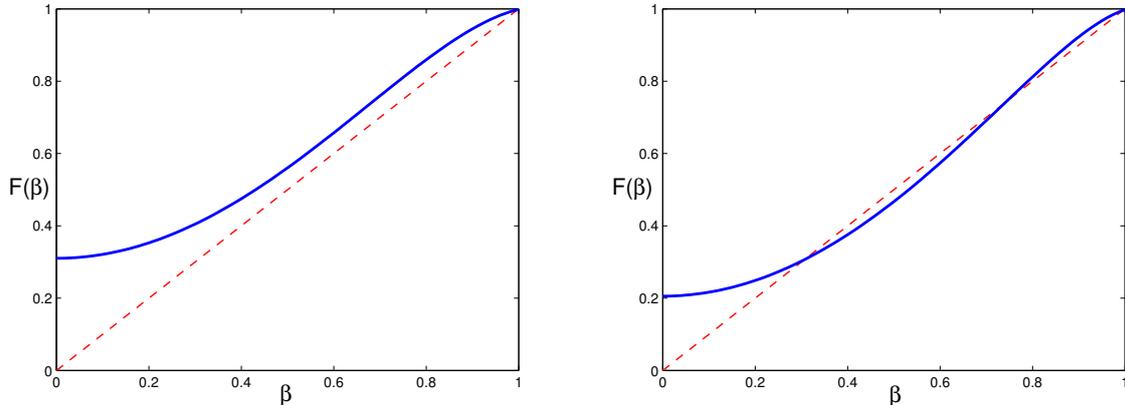


Figure 5: BLE's are fixed points satisfying $F(\beta) = \beta$, where $F(\beta)$ is the first-order autocorrelation of realized inflation in the ALM (5). The left plot illustrates a case with a unique near-unit root BLE $\beta = 0.997$. The right plot illustrates a case with three BLEs, a *low* and *high* persistence BLE $\beta^* = 0.3066$ and $\beta^* = 0.9961$ both stable under SAC-learning, and separated by an unstable BLE $\beta^* = 0.8$.

Parameters: $\delta = 0.99, \rho = 0.9, \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2} = 0.1; \gamma = 0.1$ (left plot); $\gamma = 0.075$ (right plot),

Figure 5 illustrates the existence of BLE as fixed points of the implied autocorrelation $F(\beta)$. A typical case (left plot) occurs when there is a unique near unit root BLE ($\beta^* = 0.997$). SAC-learning converges to the unique high persistence BLE. This represents an example of *persistence amplification*, that is, the persistence under learning of an AR(1) rule is much higher (often near unit root) than under RE (where the persistence is inherited from the driving variable).

A surprising result is that, even in this very simple linear univariate setting, multiple BLE may co-exist. Figure 5 (right plot) illustrates an example with three coexisting BLEs: $\beta^* = 0.3066, \beta^* = 0.8$ and $\beta^* = 0.9961$. It then depends on initial beliefs and initial states to which of these BLEs the learning process converges, as illustrated in Figure 6. Notice that the high persistent BLE exhibits persistence and volatility amplification, that is, the learning is much more volatile and much more persistent than the REE benchmark (see Figure 6 (right plot)). The SAC-learning may also switch between a low and a high

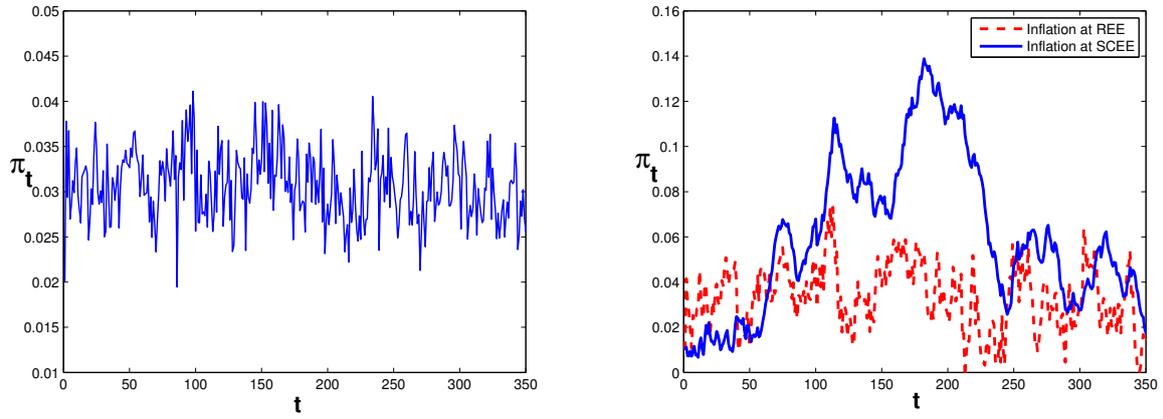


Figure 6: Convergence to low persistence BLE (left plot) or high persistence BLE (right plot) depending on initial states and initial beliefs. The high persistence BLE exhibits persistence and volatility amplification compared to the RE solution (right plot).

persistence BLE.²

²This can happen when SAC-learning is updated with a finite gain parameter.

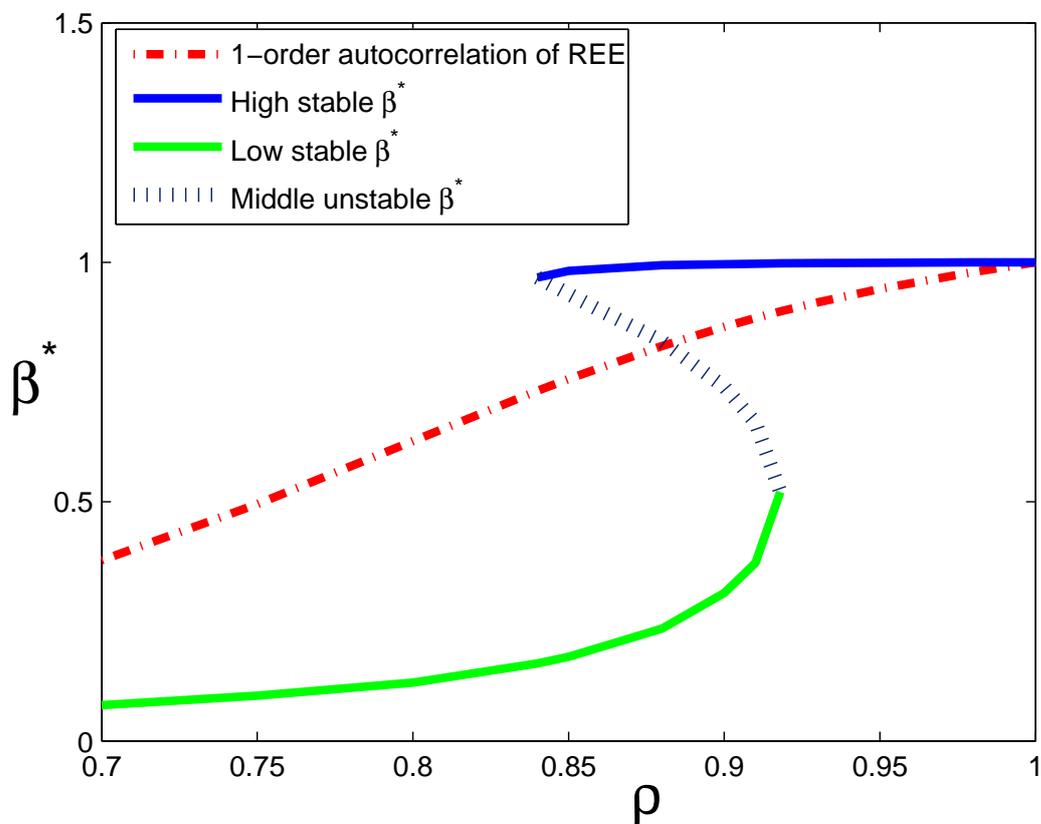


Figure 7: Critical Transitions of equilibria β^* depending on ρ where $\lambda = 0.99, \gamma = 0.075, \frac{\sigma_{\eta}^2}{\sigma_{\varepsilon}^2} = 0.1$.

Figure 7 shows the BLE curves as a function of the ρ parameter (the autocorrelation of the driving variable) and illustrates that the system exhibits multiple equilibria and critical transitions, where the system may suddenly jump between the low and the high persistence BLE.

Notice that the high persistence BLE is near unit root with a parameter $\beta^* = 0.996$ very close to 1. Notice also that for a near unit root the AR(1) rule is very close to *naive expectations*. This illustrates the point that in strong positive feedback systems such as the NKPC naive expectations may be a good forecasting rule that may be hard to beat by more complex rules.

4 Asset pricing experiments

Asset markets are another example of a strongly positive feedback systems. Optimistic asset price expectations increases demand and realized market prices. What happens in the lab, will asset prices converge to fundamentals or can (temporary) bubbles and crashes arise? In the asset pricing experiments of Hommes et al. (2005; henceforth HSTV2005) the (unknown) law of motion of the asset price is given by

$$p_t = \frac{1}{1+r} \left((1-n_t) \frac{p_{t+1,1}^e + \dots + p_{t+1,6}^e}{6} + n_t p^f + \bar{y} + \varepsilon_t \right), \quad (7)$$

where r is the risk free interest rate, \bar{y} is the mean dividend and ε_t is a small noise term. The mean dividend $\bar{y} = 3$ and the interest rate $r = 0.05$, so that subject could in principle compute the fundamental price $p^f = \bar{y}/r = 60$. Furthermore, n_t is the fraction of fundamental robot traders, who predict that the asset price will be at fundamental value $p^f = \bar{y}/r$ at all times. Robot traders are a "far from equilibrium" stabilizing force, whose fraction increases from 0 to 25%, with the distance of the market price from fundamental.

For the asset pricing model the (unknown) law of motion is thus of the form³

$$p_t = f(\bar{p}_{t+1}^e), \quad (8)$$

where \bar{p}_{t+1}^e is the average forecast of a group of six individuals⁴. The price equation has been derived from myopic mean-variance maximization. There is *strong positive feedback* as the map f is increasing and has a near unit root eigenvalue $\lambda = 1/(1+r) = 0.95$.

Figure 8 shows simulations of four benchmark expectation rules. Under rational expectations the forecast equals the fundamental price $p^f = \bar{y}/r = 60$ and the realized price falls between 59 and 61, with small random fluctuations around the fundamental price 60. Under naive expectations the price converges (almost) monotonically to the funda-

³Notice that this temporary equilibrium setup yields a 2-period ahead LtFE, since p_t is unknown yet when forecasting p_{t+1} .

⁴Most LtFEs use group sizes between 5 and 10. In a recent asset pricing LtFE Hommes et al. (2021) use group sizes of about 100, by coupling two laboratories, and show that the results are robust.

mental price after 50 periods. Under learning by sample average the price slowly increases from about 50 to 53, but convergence towards fundamentals is extremely slow. The last simulation uses a simple AR(2) heuristic

$$p_{t+1}^e = 30 + \frac{3}{2}p_{t-1} - p_{t-2}. \quad (9)$$

When estimating individual forecasting rules to our surprise for several subjects an AR(2) rule with the coefficients close to those of (9) was found. This AR(2) rule has a behavioral interpretation as an anchor and adjustment rule (Tversky and Kahnemann, 1974):

$$p_{t+1}^e = 30 + \frac{3}{2}p_{t-1} - p_{t-2} = \frac{60 + p_{t-1}}{2} + (p_{t-1} - p_{t-2}), \quad (10)$$

and this is perhaps the reason why subjects may coordinate on such a heuristic.

Figure 9 shows what happens in the lab. Three qualitatively different patterns are observed in the lab: monotonic convergence, persistent oscillations and damping oscillations. Laboratory results are thus inconsistent with rational, fundamental forecasting, although some of the groups seem to converge towards the end of the experiment. The first two groups (top panel) seem consistent with naive expectations. The persistently oscillating groups are similar to simulations with an AR(2) trend-extrapolating rule. The last two groups exhibit temporary bubbles and crashes (bottom panel).

Because of the heterogeneous results, it may be hard to fit a homogeneous expectations or learning model that fits these experimental data. Rather the experimental data suggest a heterogeneous expectations model to explain these results.

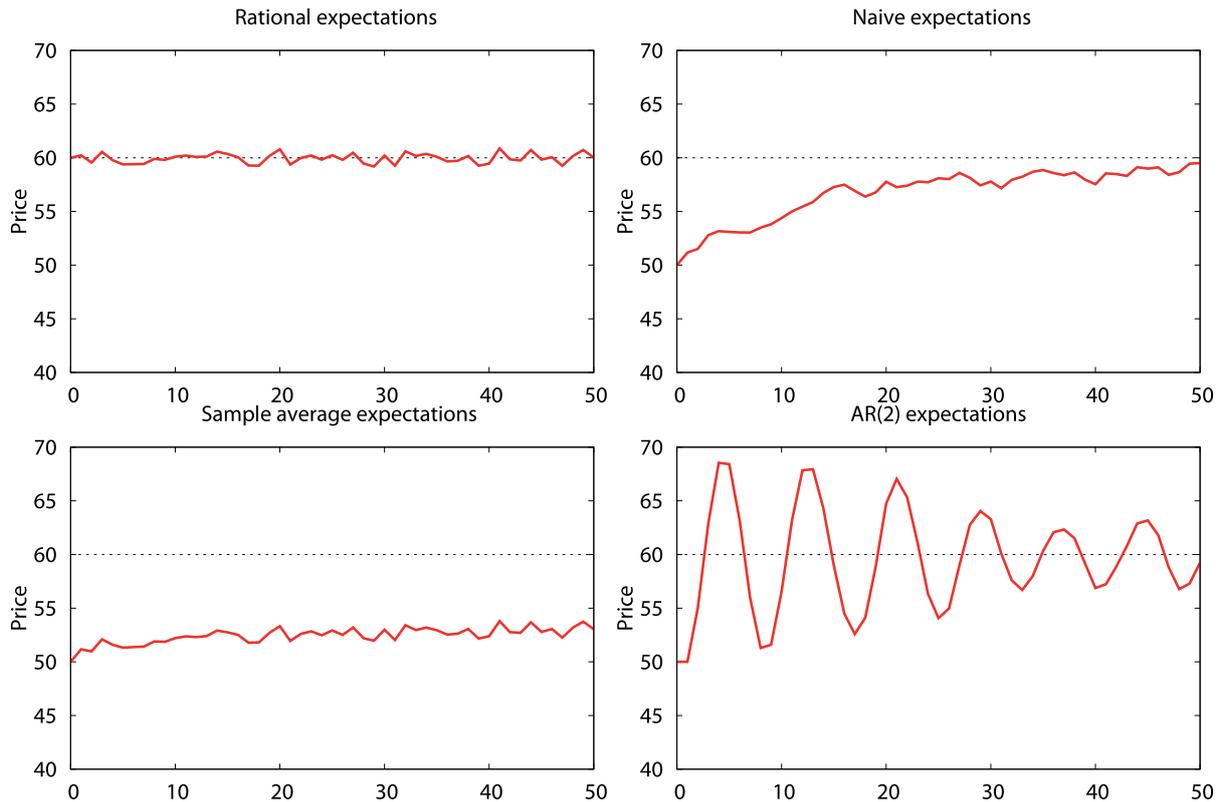


Figure 8: Simulations of four benchmark forecasting heuristics in the asset pricing experiment of HSTV2005. Under rational expectations (top left) the forecast is $p^f = \bar{y}/r = 60$ and the realized price exhibits small random fluctuations around the fundamental price 60. Under naive expectations (top right) the price converges (almost) monotonically to the fundamental price. Under learning by sample average (bottom left) the price slowly increases from about 50 to 53, but convergence is extremely slow. The last simulation uses the simple AR(2) heuristic in (10) with prices showing persistent oscillations around the fundamental .

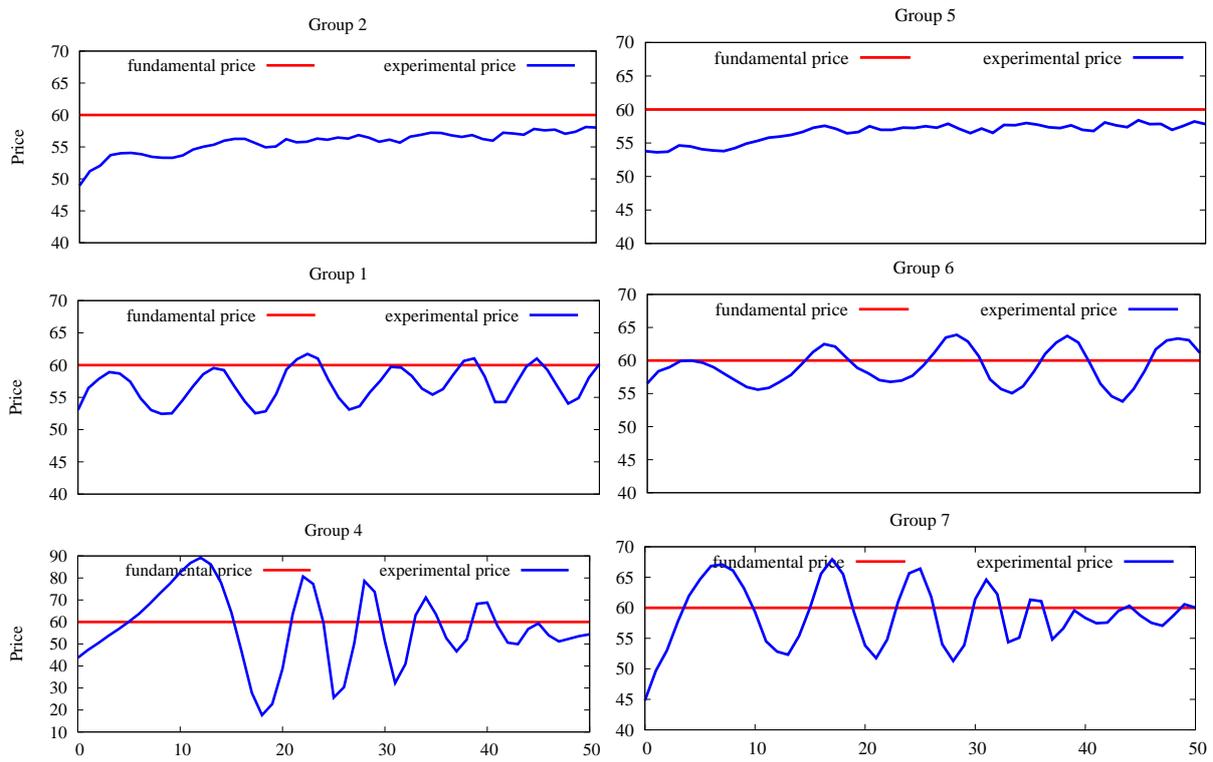


Figure 9: Six different markets from the asset market experiment HSTV2005. Three different types of aggregate behavior are observed: (i) monotonic convergence (top panels), (ii) persistent oscillations (middle panels), and (iii) bubbles and crashes that dampen towards end of experiment.

5 Heuristics switching model (HSM)

The Heuristics Switching Model (HSM), (Brock and Hommes, 1997; Anufriev and Hommes, 2012). fits these experimental data well. Agents choose from a number of simple *forecasting heuristics* based on their relative performance. Agents evaluate the *performances* of all heuristics, and tend to *switch* to more successful rules. The fractions of the belief types are gradually updated in each period and given by a discrete choice model with asynchronous updating:

$$n_{ht} = \delta n_{h,t-1} + (1 - \delta) \frac{e^{\beta U_{h,t-1}}}{Z_{t-1}},$$

where Z_{t-1} is normalization factor, so that the fractions add up to 1. Here U_{ht} is the *fitness measure* (e.g. forecasting performance, realized profits); β is the *intensity of choice*, measuring how quickly agents switch, and δ is the *asynchronous* updating parameter, representing the agents that on average stick to their strategy in a given time period.

Anufriev and Hommes, 2012, reduced the number of heuristics to four, while still fitting the experimental data well. The four heuristics are:

$$p_{1,t+1}^e = w p_{t-1} + (1 - w) p_{1,t}^e \quad ADA, w = 0.65 \quad (11)$$

$$p_{2,t+1}^e = p_{t-1} + \gamma (p_{t-1} - p_{t-2}) \quad WTR, \gamma = 0.4 \quad (12)$$

$$p_{3,t+1}^e = p_{t-1} + \gamma (p_{t-1} - p_{t-2}) \quad STR, \gamma = 1.3 \quad (13)$$

$$p_{4,t+1}^e = \frac{1}{2} (p_{t-1}^{av} + p_{t-1}) + (p_{t-1} - p_{t-2}) \quad LAA \quad (14)$$

These heuristics capture behavior that has been observed in the lab:

- *adaptive expectations* (ADA) leads to monotonically converging prices;
- *weak trend-following* rule (WTR) may lead to some overshooting, but with price oscillations converging;
- *strong trend-following* rule (STR) may lead to strong oscillations or even exploding prices;

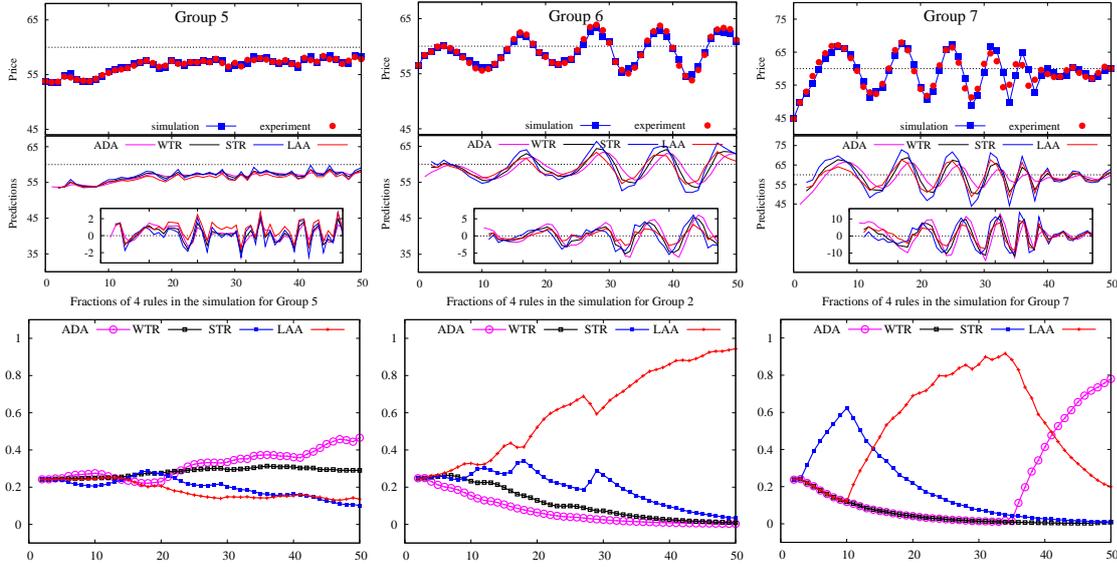


Figure 10: Heuristics switching model and experimental data. Top panels show 1-step ahead forecasts of HSM. Middle panels show forecasts and forecast errors of the four strategies. Bottom panels show the fraction of the four strategies of the HSM.

- *learning anchor and adjustment (LAA)* rule may lead to persistent oscillations⁵.

Figure 10 shows that the HSM fits the experimental data well in terms of 1-period ahead forecasting. In the case of monotonic convergence the fraction of the four forecasting heuristics are almost equal in the first half of the experiment, while adaptive expectations (ADA) dominates towards the end. In the case of persistent fluctuations the learning anchor and adjustment rule dominates, with almost 90% using the LAA rule towards the end of the experiment. Finally, the bubble and crash fluctuations are explained by strategy switching, from strong trend-following dominating the first 10-15 periods, followed by a longer phase, where the LAA rule dominates (up to more than 80%), and adaptive expectations (ADA), which had the lowest impact until period 35, taking over towards the end of the experiments with an impact of almost 80% when the price stabilizes.

⁵Note that the LAA heuristic is the only rule that may forecast a price reversal after a strong price trend, due to the time varying anchor, being an average of the last price observation and the sample average (the latter may be seen as a proxy for an equilibrium or steady state price level).

6 Positive versus negative feedback

Asset pricing experiments (Section 4) seem to be more unstable than cobweb experiments (Section 2), because subjects coordinate more easily on trend-following behavior under positive feedback. In order to test this hypothesis Heemeijer et al. (JEDC 2009) and Bao et al. (JEDC 2012) ran lab experiments, where the only difference in treatments is the sign of the feedback coefficient.

Under negative feedback (a strategic substitutes environment) the market price is given by the simple linear cobweb model

$$p_t = 60 - \frac{20}{21} \left[\sum_{h=1}^6 \frac{1}{6} p_{ht}^e \right] - 60 + \epsilon_t, \quad (15)$$

while under positive feedback (a strategic complementarity environment) the market price is given by a linear asset pricing model

$$p_t = 60 + \frac{20}{21} \left[\sum_{h=1}^6 \frac{1}{6} p_{ht}^e - 60 \right] + \epsilon_t. \quad (16)$$

Notice that (15) and (16) share the RE equilibrium price 60, and *only* differ in the sign of the slope of the linear map, -0.95 vs $+0.95$. Heemeijer et al. (2009) use small shocks ϵ_t , but here we focus on the large shock treatment of Bao et al. (2012), where unknown large permanent shocks to the fundamental steady state price occur in periods 22 and 44.

Figure 11 shows the positive and negative feedback mappings underlying the LtFE, when there are no shocks, i.e., $\epsilon_t \equiv 0$. In the negative feedback case the rational price 60 is clearly visible as a fixed point, while in the positive feedback case it is more difficult to identify the rational fixed point. Notice that under (strong) positive feedback, any price forecast is an almost self-fulfilling equilibrium. Indeed, if the eigenvalue under positive feedback were $+1$, rather than $+0.95$, the system would have a continuum of perfectly self-fulfilling rational equilibria. The question is whether the almost self-fulfilling equilibria affect the adaptive learning process and the (in)stability in the lab.

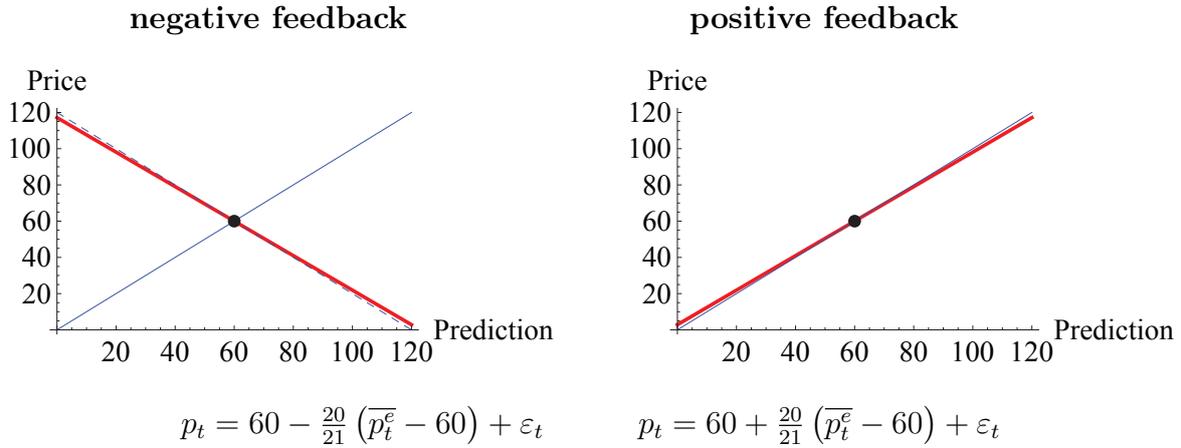


Figure 11: Negative (left) and positive (right) feedback maps with slopes -0.95 and $+0.95$. Strong positive feedback leads to almost self-fulfilling equilibria.

Positive vs Negative Feedback with Large Shocks

Bao et al. (2012) ran an experiment with large unknown permanent shocks to the fundamental price level. More precisely, there are two shocks to the fundamental equilibrium price p_t^* :

$$\begin{aligned}
 p_t^* &= 56, & 0 \leq t \leq 21, \\
 p_t^* &= 41, & 22 \leq t \leq 43, \\
 p_t^* &= 62, & 44 \leq t \leq 65.
 \end{aligned}
 \tag{17}$$

The purpose of these experiments was to investigate how the type of expectations feedback may affect the speed of learning of a new steady state equilibrium price, after a relatively large unanticipated shock to the economy.

Figure 12 shows eight negative (left panel) and eight positive (right panel) feedback groups. Aggregate behaviors under positive and negative feedback are strikingly different. Negative feedback markets tend to be rather stable, with price converging quickly to the new (unknown) equilibrium level after each unanticipated large shock. In contrast, under positive feedback prices are sluggish, converging only slowly into the direction of the fundamental value and subsequently overshooting it by large amounts.

Figure 13 shows that the HSM fits these experiments well. Under negative feedback

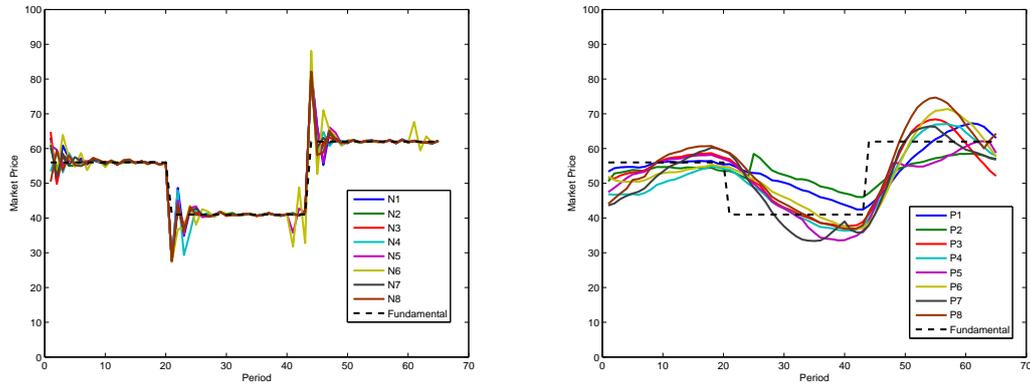


Figure 12: Eight negative (left) and positive (right) feedback groups in LtFE, with two large shocks in periods 22 and 44. After each shock under negative feedback the prices quickly converge to the new equilibrium price. Under positive feedback the market is unstable and fluctuates.

subjects coordinate on adaptive expectations or a contrarian forecasting heuristic⁶, leading to stable price behaviour. In contrast, under positive feedback subjects coordinate on trend-following heuristic amplifying price fluctuations.

⁶A contrarian rule is a trend-following rule with a negative coefficient γ , implying that a contrarian expects a price decrease after a price increase.

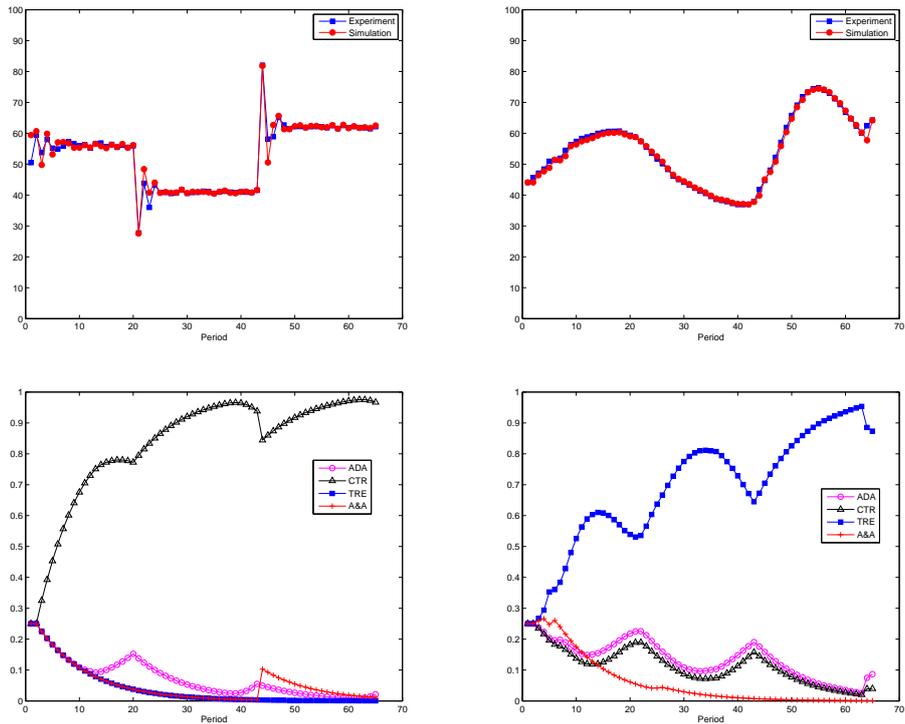


Figure 13: Heuristic Switching Model fitted to negative (left panels) and positive (right panels) feedback markets with large shocks. Under negative feedback subjects coordinate on an adaptive or a contrarian forecasting heuristic (bottom left panel), while under positive feedback coordination on trend-following heuristic (bottom right panel) amplifies price fluctuations.

7 GA-learning of a first-order heuristic

Heemeijer et al. (2009) and Bao et al. (2012) showed that simple linear *first-order heuristics* fit forecasting behavior of a majority of subjects in the lab well. A first-order forecasting heuristic only takes into account the first lag of price and forecast and the first lag of price change and, in its simplest form, is given by

$$p_t^e = \alpha p_{t-1} + (1 - \alpha)p_{t-1}^e + \beta(p_{t-1} - p_{t-2}). \quad (18)$$

The first-order heuristic is an anchor and adjustment rule, with the anchor an adaptive expectations forecast $\alpha p_{t-1} + (1 - \alpha)p_{t-1}^e$, with weight α , and $\beta(p_{t-1} - p_{t-2})$ a trend-extrapolation with coefficient β . Anufriev et al. (2019) studied genetic algorithm (GA) learning of the optimal first-order *anchor and adjustment rules* of this simple form

$$p_{i,h,t}^e = \alpha_{i,h,t} p_{t-1} + (1 - \alpha_{i,h,t}) p_{i,t-1}^e + \beta_{i,h,t} (p_{t-1} - p_{t-2}). \quad (19)$$

Here $p_{i,h,t}^e$ refers to the forecast of subject i using rule h at time t . The rule h requires two time-varying parameters: an *anchor* $\alpha_{i,h,t} \in [0, 1]$ and a *trend* $\beta_{i,h,t} \in [-1.1, 1.1]$. The heuristic (19) generalizes popular HSM heuristics: naive ($\alpha = 1, \beta = 0$), adaptive expectations ($\beta = 0$) and trend extrapolation ($\alpha = 1, \beta > 0$). The rational forecast is also nested as a special case ($\alpha = 0, \beta = 0, p_{i,t-1}^e = p^f$).

GA-learning of an optimal anchor and adjustment heuristics is implemented as follows (see Anufriev and Hommes (2019), following Arifovic (1994, 1996), and Dawid (2011)).

- Every agent i has a list of $H = 20$ different heuristics $(\alpha_{i,h}, \beta_{i,h})$.
- When agent i observes the last realized price p_{t-1} , she re-optimizes the rules with GA evolutionary operators:

1. *reproduction; (survival of the fittest)*: sample (with replacement) 20 new heuristics from the old depending on their *hypothetical forecasting performance*;

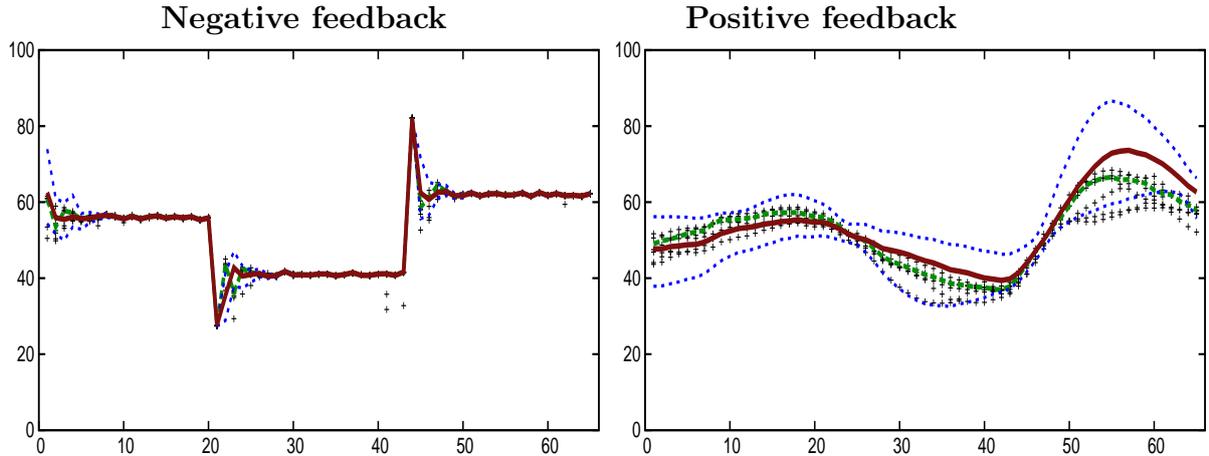


Figure 14: 65-period ahead Monte Carlo simulations (1000); experimental data Bao et al. (2012). Green dashed line and black pluses show the experimental median and group observations, respectively. Red thick line is the median simulation and blue dotted lines are the 95% confidence interval for the GA model.

2. *mutation*: with some small probability “mutate” them (i.e., modify $(\alpha_{i,h}, \beta_{i,h})$ of each heuristic);
3. *election*: compare the new and the old heuristics in terms of their hypothetical forecasting performance – pick the better ones.

The GA process mimics *natural selection*: worse forecasting heuristics are likely to be *replaced* by better ones, with some small random mutation; *inefficient experimentation* is screened out by the election operator.

Figure 14 shows 65-period ahead Monte Carlo simulations (1000). The GA learning model fits the experimental data remarkably well.

Figure 15 illustrates which heuristics were learned by the GA agents in the lab. The figure shows the median and the mean (with 95% and 90% CI) for 1000 runs in a Monte Carlo simulation of the anchor or price weight α_t and the trend extrapolation coefficient β_t , which were selected by the six GA agents for both negative and positive feedback.

Under *negative feedback* agents learn to use *adaptive expectations*, with an optimal

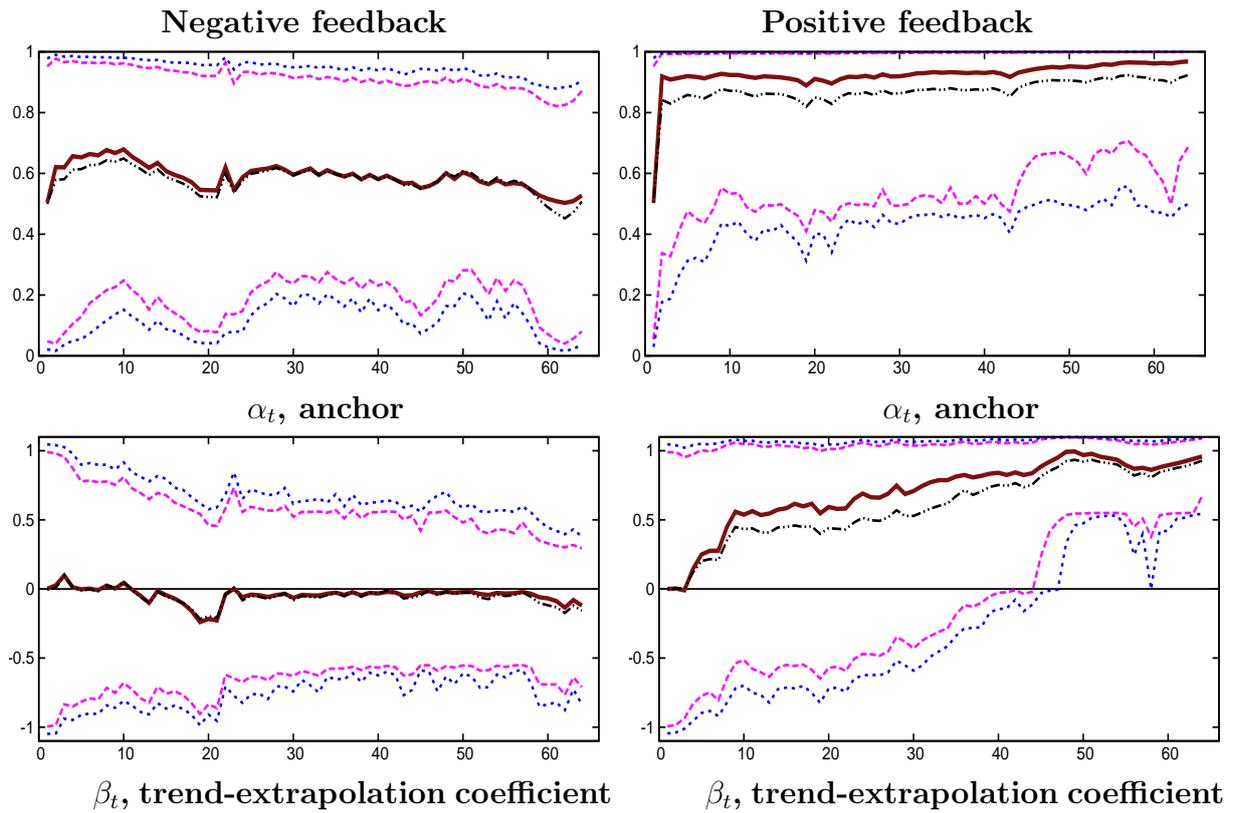


Figure 15: GA-learning of optimal anchor α_t (top) and trend β_t (bottom) parameters for negative (left) and positive (right) feedback treatments (experimental data Bao et al. (2012)). Under negative feedback the heuristic converges to adaptive expectations, while under positive feedback GA-learning converges to trend-extrapolation. The wide CIs show that the forecasts are very heterogeneous, consistent with estimated rules.

heuristic

$$p_{i,t}^e \approx 0.5p_{t-1} + 0.5p_{i,t-1}^e.$$

Under *positive feedback* agents learn to become *trend-follower*, with an optimal heuristic:

$$p_{i,t}^e \approx 0.95p_{t-1} + 0.05p_{i,t-1}^e + 0.9(p_{t-1} - p_{t-2}).$$

Negative feedback systems thus favor stabilizing adaptive expectations, while in positive feedback systems agents tend to coordinate on trend-extrapolating heuristics, with an anchor close to the last observed price, amplifying fluctuations.

8 Managing trend-following behavior through monetary policy

Trend-following strategies are destabilizing in positive feedback systems and may lead to persistent deviations from fundamentals. How can policy manage coordination on trend-following heuristics and stabilize aggregate behavior of the system?

Consider this question within the simple three-equations New Keynesian model, with output, inflation and an inflation targeting monetary policy interest rate rule:

$$\begin{aligned}
 y_t &= y_{t+1}^e - \varphi(i_t - \pi_{t+1}^e) + \epsilon_t, & \text{output} \\
 \pi_t &= \lambda y_t + \rho \pi_{t+1}^e + v_t, & \text{inflation} \\
 i_t &= \text{Max}\{\bar{\pi} + \phi_\pi(\pi_t - \bar{\pi}), 0\} & \text{monetary policy rule}
 \end{aligned}$$

In this economy there are two forecasts to be made: π_{t+1}^e for inflation and y_{t+1}^e for output gap. The system has different types of expectational feedback. There is *positive feedback* from expectations π_{t+1}^e and y_{t+1}^e on output and inflation. However, there is *negative feedback* from the inflation targeting interest rate rule on the output gap. By increasing the magnitude of the inflation targeting policy parameter ϕ_π , the Central Bank adds negative feedback to the macro-economy. Can a more aggressive inflation targeting rule stabilize inflation and output? Assenza et al. (2021) conducted a LtFE lab experiment to address this question with agents forming incentivised expectations in the lab.

Figure 16 shows the results in the lab for different values of the policy parameter ϕ_π . For weak inflation targeting ($\phi_\pi = 1$) the NK expectations feedback system has an eigenvalue $\lambda = 1$ and the experimental economy is highly unstable, with exploding inflation and output in most groups. For $\phi_\pi = 1.005$, slightly larger than 1, the expectations feedback system has eigenvalues inside the unit circle and the REE is determinate. In the lab, however, the economy is still highly unstable, with most groups diverging. For a slightly larger value $\phi_\pi = 1.015$, expectations feedback from the interest rate towards the output

gap becomes negative and, as a result, the NK economies in the lab are much more stable, with small fluctuations around the inflation target. Finally, for $\phi_\pi = 1.5$, the more aggressive monetary policy adds sufficiently strong negative feedback to fully stabilize the NK economies.

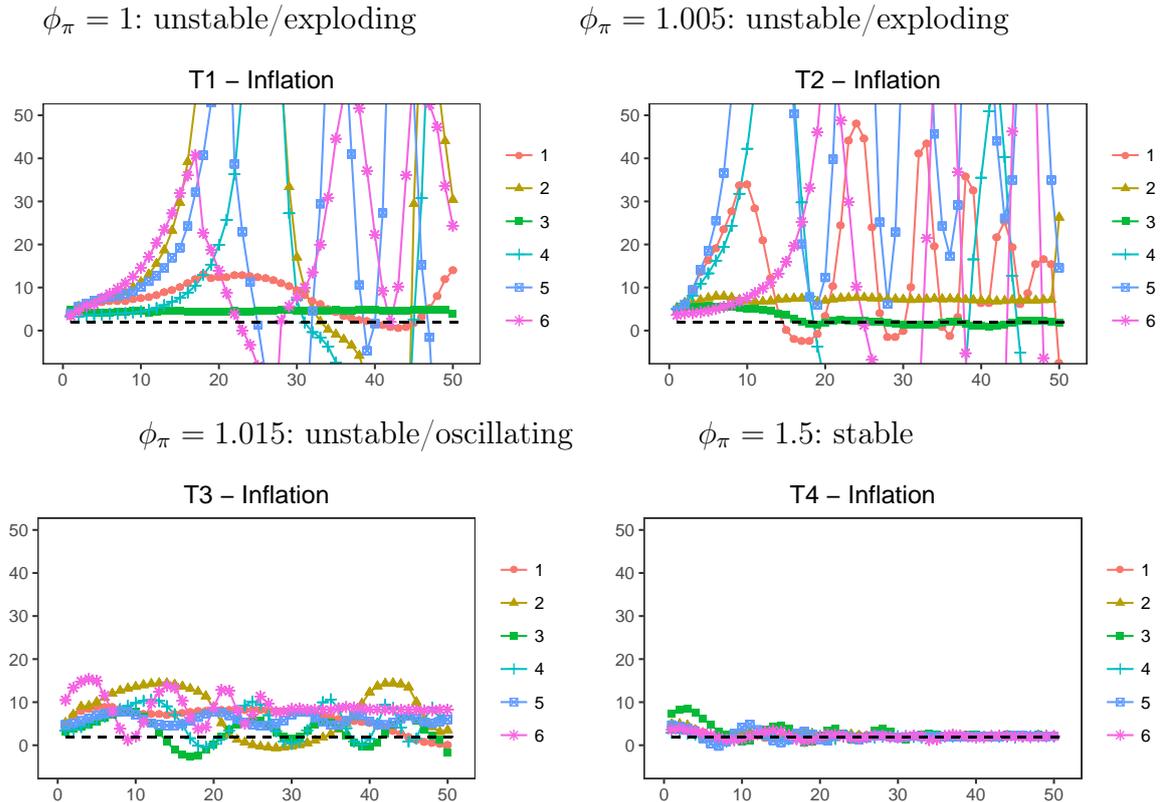


Figure 16: A more aggressive Taylor rule, that is, a higher policy coefficient ϕ_π , adds negative feedback to the NK system and stabilizes inflation (and output) when the Taylor rule is sufficiently aggressive.

The HSM model explains these differences in the experimental treatments, as illustrated in Figures 17–19. Under weak inflation targeting ($\phi_\pi = 1$ or $\phi_\pi = 1.005$; Figure 17), subjects coordinate on a strong trend-following forecasting rule and inflation is highly unstable and explodes. Under medium inflation targeting ($\phi_\pi = 1.015$; Figure 18), subjects coordinate on a learning anchor and adjustment (LAA) rule and inflation and output gap fluctuate around their target levels. Under strong inflation targeting ($\phi_\pi = 1.5$; Fig-

ure 19), subjects coordinate on adaptive expectations (AE) and inflation and output gap stabilize.

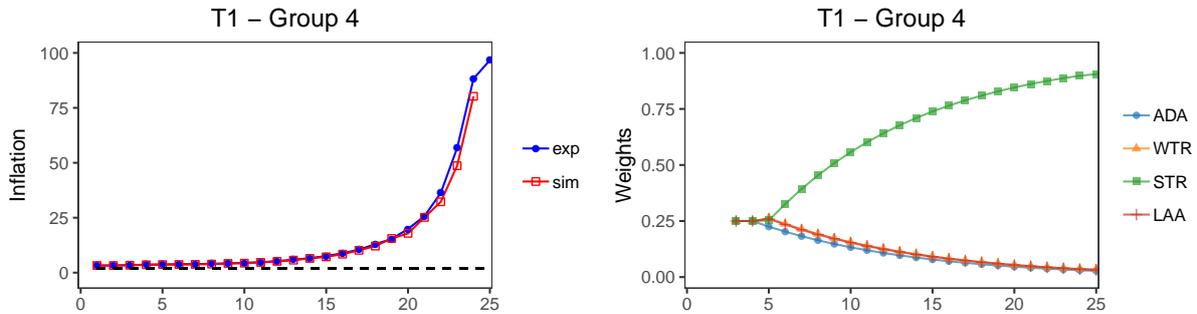


Figure 17: Under weak inflation targeting ($\phi_\pi = 1$ or $\phi_\pi = 1.005$), agents coordinate on a strong trend-extrapolating rule (right plot) and the economy is highly unstable with inflation (and output gap) exploding (left plot).

This simple NK model with heterogeneous heuristics shows that the micro and macro behavior depend crucially on policy parameters. The Central Bank therefore can influence the micro and macro behavior of the system by monetary policy. The policy parameter affects the adaptive learning process, in particular, policy affects the performance and impact of different heuristics and therefore affects the choice of the forecasting heuristics. A more aggressive inflation targeting policy weakens the positive feedback and therefore makes it more unlikely that agents will coordinate on a destabilizing trend-following heuristic.

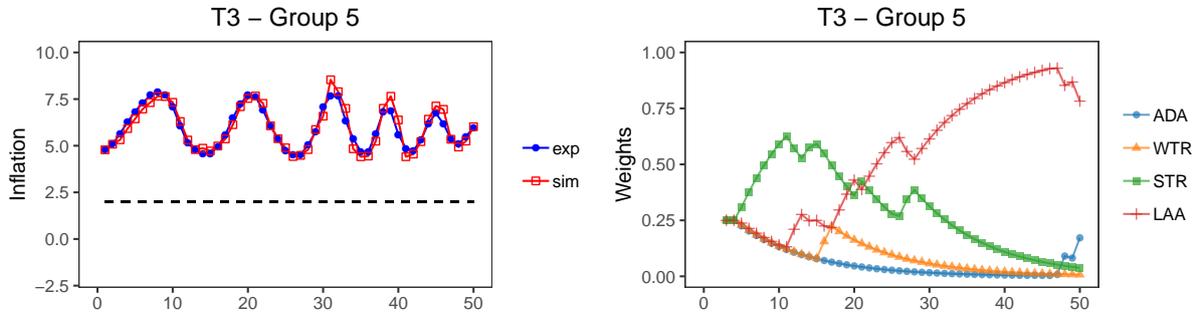


Figure 18: Under intermediate inflation targeting ($\phi_\pi = 1.015$), with small negative FB from inflation targeting, agents coordinate on the learning-anchor and adjustment (LAA) rule (right plot) and inflation (and output gap) fluctuate around their target levels (left plot).

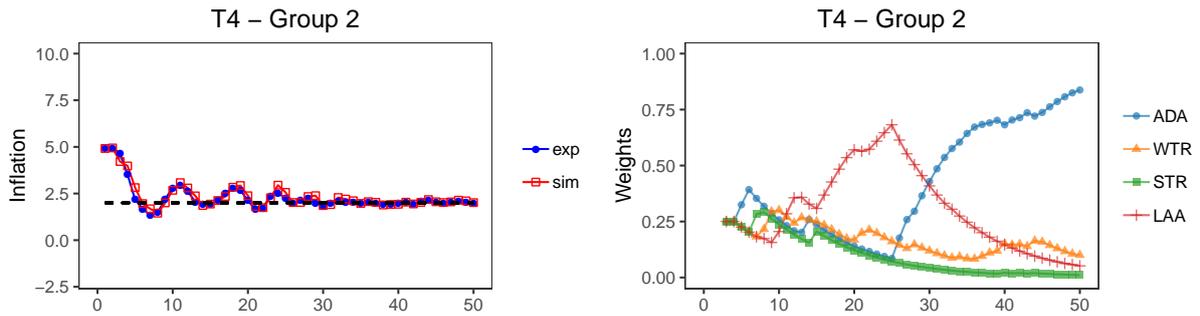


Figure 19: Under strong inflation targeting ($\phi_\pi = 1.5$) agents coordinate on adaptive expectations (right plot) and inflation (and output gap) stabilize

9 Adaptive behavior in the lab

Lucas (1986) viewed the stability under learning as an *experimentally testable hypothesis*. This has triggered a large literature on LtFEs (see the pioneering work of Marimon, Spear and Sunder (1993) and Marimon and Sunder (1993); see Hommes (2021) for a review. .

Hennequin et al. (2025) use learning-to-forecast experiments (LtFEs) to study the stability of expectations in the lab in a multivariate setting. The (in)stability in the lab is characterized by the leading eigenvalues of the expectations feedback system. Surprisingly, this empirical characterization from the lab, is different from what most learning theories would predict. Learning theory usually predicts that eigenvalues inside the unit circle lead to stable behavior. The heatmap in Figure 20 shows, however, that in the laboratory unstable behavior and persistent deviations from fundamentals, arise for eigenvalues *inside* the unit circle, consistent with experimental results from the asset pricing experiments (see Section 4) and the New Keynesian experiments (see Section 8). Using a behavioral mixed heuristics switching model (HSM) calibrated to the LtFE, Hennequin et al. (2025) find an instability region inside the unit circle, with real or complex near unit root eigenvalues. Such experimental economies are unstable, because subjects in the lab coordinate on trend-following behavior amplifying temporary bubbles and crashes.

This suggests a general role for policy in stabilizing an economy with expectations feedback. Policy should add negative feedback to the system and move policy parameters so that the eigenvalues are shifted into the blue stable region in the heatmap, away from the purple instability region. The general mechanism behind these stabilizing policies is that weakening positive feedback makes coordination on trend-following heuristic less likely and coordination on stabilizing adaptive expectations more likely. .

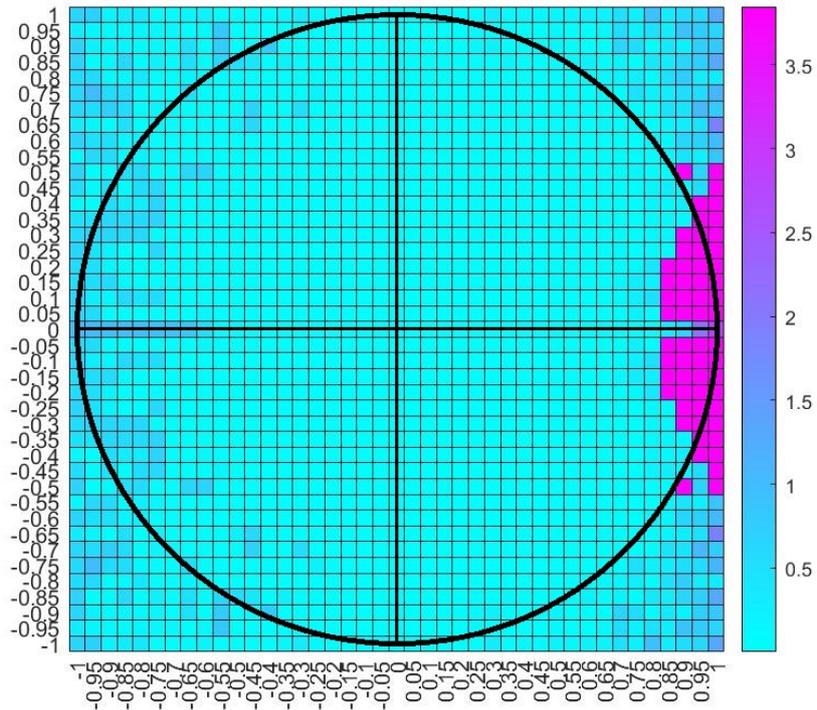


Figure 20: Heatmap showing the relative absolute deviations (RAD) simulations using the mixed HSM for a grid of eigenvalues in the complex unit square. The x-axis shows the real part, the y-axis the complex part of the (leading) eigenvalues. The unit circle is shown for clarity, illustrating the stability boundary of typical learning theories. The purple region is our predicted *instability region* in the lab, which partly lies inside the unit circle.

10 Concluding Remarks

We have discussed the use of simple adaptive forecasting heuristics as a model for more realistic expectations in macroeconomics. Some behavioral takeaways are:

- Simple *heuristics using limited information* may sometimes converge to FIRE; in such an environment coordination on the rational outcome may arise through coordination on a limited information heuristic.
- Naive expectations performs rather poorly in the cobweb model environment with negative feedback, with large and systematic forecasting errors.
- In contrast, naive expectations performs rather well in a near-unit root setting, such as an asset pricing or New Keynesian model. In particular for a random walk the naive forecast is the optimal forecast.
- Adaptive learning of an (optimal) AR1 (or AR2) heuristic may lead to limited information misspecification equilibria, characterized by almost self-fulfilling near unit root behavior and multiple equilibria;
- Positive feedback may lead to persistent deviations from fundamentals and generate instability and temporary bubbles and crashes in the laboratory, with a group of human subjects learning to coordinate on trend-following or anchor and adjustment heuristics.
- Heuristics switching based on relative performance fits laboratory and empirical macro data well.
- Learning simple heuristics may work well in highly complex macro environment and may be hard to beat by more complex forecasting rules. ;

- Negative feedback policies may shift beliefs away from trend-following to adaptive expectations and stabilize strongly positive feedback macro systems.

In general, heuristics may work as a coordination device of an adaptive learning process. However, the adaptive learning process may exhibit path-dependence either enforcing stability (e.g., through adaptive expectations) or leading to persistent deviations from fundamentals by enforcing unstable bubbles (through trend-following rules) and crashes. Policies adding negative feedback can manage expectations and shift beliefs from unstable trend-following to stabilizing adaptive expectations.

References

1. Adam, K. (2007). Experimental evidence on the persistence of output and inflation. *The Economic Journal*, 117 (520), 603–636.
2. Angeletos, George-Marios, and Chen Lian. 2016. “Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination.” In: *Handbook of Macroeconomics*. Vol. 2A, edited by John B. Taylor and Harald Uhlig, 1065–1240.
3. Anufriev, M. and Hommes, C. (2012). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(4):35–64.
4. Anufriev, Mikhail, Cars Hommes, and Tomasz Makarewicz, (2019). “Simple Forecasting Heuristics that Make Us Smart: Evidence from Different Market Experiments.” *Journal of the European Economic Association* 17 (5): 1538–84.
5. Arifovic, Jasmina. 1994. “Genetic Algorithm Learning and the Cobweb Model.” *Journal of Economic Dynamics and Control* 18 (1): 3–28.
6. Arifovic, J., 1996. The behavior of the exchange rate in the genetic algorithm and experimental economies. *Journal of Political Economy* 104 (3), 510–541.
7. Arifovic, J., Hommes, C.H. and Salle, I.L., (2019), Learning to believe in simple equilibria in a complex OLG economy, *Journal of Economic Theory* 183, 106-182.
8. Assenza, T., P. Heemeijer, C. H. Hommes, and D. Massaro. 2021. “Managing Self-Organization of Expectations through Monetary Policy: A Macro Experiment.” *Journal of Monetary Economics* 117, 170–186.
9. Bao, Te, Cars Hommes, Joep Sonnemans, and Jan Tuinstra. (2012). “Individual Expectations, Limited Rationality and Aggregate Outcomes.” *Journal of Economic*

Dynamics and Control 36 (8): 1101–20.

10. Branch, William A. 2006. “Restricted Perceptions Equilibria and Learning in Macroeconomics.” In *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*, edited by David Colander, 135–60. New York: Cambridge University Press.
11. Brock, William A., and Cars H. Hommes. 1997. “A Rational Route to Randomness.” *Econometrica* 65 (5): 1059–95.
12. Bullard, J. (1994), ‘Learning equilibria’, *Journal of Economic Theory* 64, 468-485.
13. Dawid, H. (2011), *Adaptive learning by genetic algorithms: Analytical results and applications to economic models*, Springer Science & Business Media.
14. Erev, I. and Roth, A.E. (1998), Predicting how people play games: reinforcement learning in experimental games with unique mixed strategy equilibria, *American Economic Review* 88, 841-81.
15. Evans, G. and Honkapohja, S., (2001), *Learning in macroeconomics*.
16. Gigerenzer, G., Todd, P. M., & The ABC Research Group. (1999). *Simple heuristics that make us smart*. Oxford University Press.
17. Hart, Sergiu, (2005), Adaptive heuristics, *Econometrica*, Vol. 73, No. 5, 1401-1430.
18. Heemeijer, Peter, Cars Hommes, Joep Sonnemans, and Jan Tuinstra. 2009. “Price Stability and Volatility in Markets with Positive and Negative Expectations Feedback: An Experimental Investigation.” *Journal of Economic Dynamics and Control* 33 (5): 1052–72.
19. Hennequin, M., Hommes, C.H., Kopanyi-Peuker, A. and Levelt, E. (2025) Adaptive behavior in the lab, Working Paper University of Amsterdam.

20. Hommes, C.H., Lux, T., (2013), Individual expectations and aggregate behavior in learning to forecast experiments, *Macroeconomic Dynamics* 17(2), March 2013.
21. Hommes, C.H., Kopányi-Peuker, A. and Sonnemans, J. (2021), Bubbles, crashes and information contagion in large-group asset market experiments, *Experimental Economics* 24, 414-433
22. Hommes, Cars, Joep Sonnemans, Jan Tuinstra, and Henk van de Velden. 2005. “Coordination of Expectations in Asset Pricing Experiments.” *Review of Financial Studies* 18 (3): 955–80.
23. Hommes, C.H., Sonnemans, J., Tuinstra, J. and Velden, H. van de, (2007), Learning in cobweb experiments, *Macroeconomic Dynamics* 11, S1, 8-33.
24. Hommes, Cars, and Gerhard Sorger. (1998). “Consistent Expectations Equilibria.” *Macroeconomic Dynamics* 2 (3): 287–321.
25. Hommes, Cars, and Mei Zhu. 2014. “Behavioral Learning Equilibria.” *Journal of Economic Theory* 150: 778–814.
26. Lansing, Kevin J. 2009. “Time-Varying U.S. Inflation Dynamics and the New Keynesian Phillips Curve.” *Review of Economic Dynamics* 12 (2): 304–26.
27. Lucas, Robert E., Jr. 1986. “Adaptive Behavior and Economic Theory.” *Journal of Business* 59 (4, part 2) S401–26
28. Marimon, R., Spear, S. E., and Sunder, S. (1993). Expectationally driven market volatility: An experimental study. *Journal of Economic Theory*, 61(1):74–103.
29. Marimon, R. and Sunder, S. (1993). Indeterminacy of equilibria in a hyperinflationary world: Experimental evidence. *Econometrica*, 61:1073–1107.

30. Mauersberger, F. and Nagel, R. (2018). Levels of reasoning in keynesian beauty contests: A generative framework. In Hommes, C. and LeBaron, B., editors, *Heterogeneous agents models*, volume 4 of *Handbook of Computational Economics*, pages 541–634. Elsevier.
31. Benjamin Moll, *The Trouble with Rational Expectations in Heterogeneous Agent Models: A Challenge for Macroeconomics*, Working Paper November 2024
32. Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, pages 315–335.
33. Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326.
34. Sargent, T. (1993), *Bounded rationality in macroeconomics*, Oxford University Press.
35. Simon, H.A. (1955), A Behavioral Model of Rational Choice, *The Quarterly Journal of Economics*, 69(1), 99-118.
36. Tversky, Amos and Daniel Kahneman, “Judgment under Uncertainty: Heuristics and Biases,” *Science*, 1974, 185 (4157), 1124–1131.
37. Michael Weber, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny and Jorge Ponce (2025), Tell me something I don’t already know: learning in low- and high-inflation settings, *Econometrica* Vol. 93, Issue 1 (January 2025).
38. Young, Peyton, 2008, Adaptive heuristics, In: *New Palgrave Dictionary of Economics*, Second Edition, Blume, L. and Durlauf, S. (Eds.)