

Behavioral and Experimental Macroeconomics and Policy Analysis: A Complex Systems Approach[†]

CARS HOMMES*

This survey discusses behavioral and experimental macroeconomics, emphasizing a complex systems perspective. The economy consists of boundedly rational heterogeneous agents who do not fully understand their complex environment and use simple decision heuristics. Central to our survey is the question of under which conditions a complex macro-system of interacting agents may or may not coordinate on the rational equilibrium outcome. A general finding is that under positive expectations feedback (strategic complementarity)—where optimistic (pessimistic) expectations can cause a boom (bust)—coordination failures are quite common. The economy is then rather unstable, and persistent aggregate fluctuations arise strongly amplified by coordination on trend-following behavior leading to (almost-)self-fulfilling equilibria. Heterogeneous expectations and heuristics switching models match this observed micro and macro behavior surprisingly well. We also discuss policy implications of this coordination failure on the perfectly rational aggregate outcome and how policy can help to manage the self-organization process of a complex economic system. (JEL C63, C90, D91, E12, E71, G12)

1. Introduction

The financial crisis of 2007–08 and the subsequent Great Recession, the most severe economic crisis since the Great Depression in the 1930s, have increased

concerns from policy makers and academics about the empirical relevance of the standard representative rational agent framework in macroeconomics. In an often-quoted speech during the crisis in November 2010, European Central Bank then-Governor

*CeNDEF, University of Amsterdam, and Tinbergen Institute; European Central Bank (ECB), Frankfurt; Bank of Canada. This survey has been written during my stay as Duisenberg Research Fellow at the ECB, Frankfurt, December 2017–March 2018. This research reflects my subjective views, not those of the ECB and/or the Bank of Canada. I am grateful for the financial and intellectual support from the ECB, in

particular to Luc Laeven and Frank Hartmann for supporting this research. Detailed comments of four referees and the Editor greatly improved this survey. I also would like to thank Matthias Weber for detailed comment on an earlier draft.

[†] Go to <https://doi.org/10.1257/jel.20191434> to visit the article page and view author disclosure statement(s).

Jean-Claude Trichet (2010) expressed these concerns as follows:

When the crisis came, the serious limitations of existing economic and financial models immediately became apparent. Macro models failed to predict the crisis and seemed incapable of explaining what was happening to the economy in a convincing manner. As a policy-maker during the crisis, I found the available models of limited help. In fact, I would go further: in the face of the crisis, we felt abandoned by conventional tools.

Macroeconomists have raised similar concerns. For example, Blanchard (2014) stressed that

The main lesson of the crisis is that we were much closer to “dark corners”—situations in which the economy could badly malfunction—than we thought. Now that we are more aware of nonlinearities and the dangers they pose, we should explore them further theoretically and empirically ... If macroeconomic policy and financial regulation are set in such a way as to maintain a healthy distance from dark corners, then our models that portray normal times may still be largely appropriate. Another class of economic models, aimed at measuring systemic risk, can be used to give warning signals that we are getting too close to dark corners, and that steps must be taken to reduce risk and increase distance.

The most important class of macro-models, before the crisis commonly used by central banks and other policy institutions, are the dynamic stochastic general equilibrium (DSGE) models. In response to the critique above, since the crisis DSGE macro-models have been adapted and extended by including financial frictions within the new Keynesian (NK) framework, for example, in Cúrdia and Woodford (2010, 2016); Gertler and Karadi (2011, 2013); Christiano, Motto, and Rostagno (2010); and Gilchrist, Ortiz, and Zakrajšek (2009). These extensions, however, maintain the standard rationality framework of mainstream macroeconomics assuming infinite horizon utility and profit

maximization and fully rational expectations. In the speech quoted above, Trichet (2010) went much further:

The atomistic, optimising agents underlying existing models do not capture behaviour during a crisis period. We need to deal better with heterogeneity across agents and the interaction among those heterogeneous agents. We need to entertain alternative motivations for economic choices. Behavioral economics draws on psychology to explain decisions made in crisis circumstances. Agent-based modelling dispenses with the optimisation assumption and allows for more complex interactions between agents.

Since the outbreak of the financial-economic crisis a heavy debate among macroeconomists about the future of macroeconomic theory has emerged. The recent special issue on “Rebuilding Macroeconomic Theory” in the *Oxford Review of Economic Policy* (2018, volume 34, issues 1–2) collects a number of recent discussions on this topic. Stiglitz (2018) is particularly critical of DSGE models; Christiano, Eichenbaum, and Trabandt (2018) provide a detailed reply defending the DSGE approach. Based on questionnaires and two conferences, Vines and Wills (2018) conclude that four main changes to the core model in macroeconomics are recommended: (i) to emphasize financial frictions, (ii) to place a limit on the operation of rational expectations, (iii) to include heterogeneous agents, and (iv) to devise more appropriate microfoundations. There have also been more radical proposals for changing macro by a paradigm shift to using an interdisciplinary complex systems approach, behavioral agent-based models, and simulation (rather than analytical tools), for example, Battiston et al. (2016), Bookstaber and Kirman (2018), Haldane and Turrell (2018), and Dawid and Delli Gatti (2018).

This paper surveys some of the literature taking such a more radical, behavioral departure from the standard representative rational agent model emphasizing the role

of *nonrational expectations and bounded rationality in stylized complexity models*. There is a large behavioral macroeconomics literature on this topic, but many mainstream macroeconomists seem to be largely unaware of it. We argue that allowing for learning and heterogeneous expectations enriches the standard models with nonlinearities and many empirically relevant features, such as boom and bust cycles. In the last two decades a rich behavioral theory of expectations that fits empirical time series observations, laboratory experiments, and survey data has emerged that should become part of the standard toolbox for policy analysis.¹

Behavioral economics has become widely accepted and, one could argue, belongs to the mainstream at least since the Nobel Prizes of George Akerlof in 2001 and Daniel Kahnemann in 2002. But much of the research in the area of behavioral economics focused on individual behavior and macroeconomists, until recently, have argued that behavioral biases wash out at the aggregate level. Behavioral finance has also become well established and, for example, much of the work of the 2013 and 2017 Nobel Prize winners Robert Shiller and Richard Thaler fits into behavioral finance. Recently, however, macroeconomists show an increased interest in behavioral modeling. For example, at the NBER summer institute Andrew Caplin and Mike Woodford have organized workshops on behavioral macroeconomics since 2015 and a *JEL* code (E03) for behavioral macroeconomics has existed since 2017. Experimental

economics is also a well-established field and, one could argue, has become part of the mainstream since the Nobel Prizes of Reinhard Selten in 1994 and Vernon Smith in 2002. Most lab experiments, however, focused on individual decision making or on strategic interactions in games with two or three players. Although market experiments with small groups (say six to ten subjects) go back a long way, to at least the double auction experiments of Smith (1962) and the influential asset market bubble experiments of Smith, Suchanek, and Williams (1988), most macroeconomists have ignored laboratory experiments as a research method. But macroeconomics could benefit from lab experiments in a similar way as microeconomics has done, and macroeconomists should address the question: if a macro theory does not work in a simple controlled laboratory environment, why would it work in reality? Experimental macroeconomics is becoming increasingly popular as a complementary method to studying stylized macrosystems and falsifying macro theory in controlled laboratory environments; see, for example, the collection of papers in Duffy (2014) and the recent handbook chapters Duffy (2016), Arifovic and Duffy (2018), and Mauersberger and Nagel (2018).

The starting point of our survey is the development of theories of learning in macroeconomics originating more than 30 years ago, when macroeconomists became aware of the multiplicity of (rational) equilibria in standard macro-model settings. As a direct motivation and inspiration for this survey we use the following quote from Lucas (1986) concerning stability or learning theory [emphasis added]:

Recent theoretical work is making it increasingly clear that the *multiplicity of equilibria* ... can arise in a wide variety of situations involving sequential trading, in competitive as well as finite agent games. All but a few of these equilibria are, I believe, behaviorally uninteresting; They do not describe *behavior that collections*

¹In a related but different survey Woodford (2014) discusses the role of nonrational expectations within the new Keynesian modeling framework. While Woodford restricts attention to homogeneous expectations and stresses close-to-rational expectations, such as near-rational expectations (Woodford 2010, Adam and Woodford 2012) and rational belief equilibria (Kurz 1997), we will stress behavioral features and parsimonious forecasting heuristics and emphasize the role of heterogeneous expectations.

of *adaptively behaving people* would ever hit on. I think an appropriate *stability theory* can be useful in weeding out these uninteresting equilibria ... 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.

A key question for macroeconomic behavior then is: *what is the aggregate behavior that a collection of adaptively behaving individuals will learn to coordinate on?*² A second key question is: *how can policy affect this complex coordination process?*³ To discuss these questions and survey the state of the art of the literature two topics are of particular interest and deserve a brief discussion in this introduction: (i) complex systems and (ii) macro laboratory experiments.

There is no universal definition of a complex system, but there are two important characteristics that we will stress:² (i) nonlinearity and (ii) heterogeneity. Nonlinearities can lead to *multiple equilibria* and, as a consequence, small changes at the micro level may amplify and lead to critical transitions or tipping points at the macro level. Figure 1 illustrates the phenomenon of a critical transition—see Scheffer (2009) for an extensive discussion. When nonlinearities are mild, a change in parameters only causes a gradual change in the unique stable steady state of the system. When nonlinearities become stronger, then a small change in parameters may lead to a larger change in the stable steady state of the system, but the change is

still continuous and reversible. In the presence of very strong nonlinearities, *multiple* steady states arise and catastrophic changes from a “good” steady state to a “bad” or “crisis” steady state of the system may occur after small changes of parameters (e.g., Scheffer 2009, Scheffer et al. 2012). After such a catastrophic change, the system can not easily be recovered and pushed back to the “good” steady state (see the caption of figure 1). Such strong nonlinearities can model the “dark corners” of the economy Blanchard (2014) refers to. It is very important to understand the key nonlinearities of the economy, in order to control policy parameters to prevent the system from undesirable critical transitions and sudden collapse. Standard DSGE models have been criticized for not being able to predict the financial-economic crisis. Such a critique may be unfair, because crises in complex systems are very hard to predict. However, what has been more critical for the standard DSGE model is its almost entire focus on (log) linearized models with fully rational agents and a unique equilibrium. In such models, by assumption, a crisis through a critical transition can never exist. A realistic model of the macroeconomy should allow for the *possibility* of a crisis other than through large exogenous shocks.

A second important aspect of complex systems is that they consist of multiple (often many) *heterogeneous agents*, who interact with each other. A multi-agent complex macro-system can not be reduced to a single, individual agent system, but its interactions at the micro level must be studied to explain its aggregate behavior.³ Complex systems exhibit *emergent macro behavior* as the aggregate outcome of micro interactions.

²Another important aspect of complex systems that is receiving much attention in recent work concerns networks. For example, financial networks may have increased systemic risk and may have caused cascades that have exaggerated the global financial-economic crisis. This aspect of complex systems will not be dealt with here. The interested reader is, for example, referred to Iori and Mantegna (2018) and Goyal (2018).

³The key observation that macro behavior in a complex system can not be reduced to micro behavior has been nicely summarized in the title of one of the first and seminal papers on complexity: “More Is Different,” by Anderson (1972).




Figure 1. Multiple Steady States and Critical Transitions or Tipping Points

Notes: Panel A: when nonlinearities are mild, the steady state is unique and a change in parameters only leads to a gradual change in the stable steady state of the system. Panel B: when nonlinearities become stronger, a small change in parameters may lead to a larger change in the stable steady state of the system, but the change is still continuous and reversible. Panel C: with very strong nonlinearities, multiple steady states coexist and catastrophic changes from a “good” steady state to a “bad” or “crisis” steady state of the system may occur after small changes of a parameter. At the point F_1 a catastrophic change occurs and the system jumps from the “good” upper stable steady state to the “bad” lower steady state. After such a catastrophic change, the system can not easily be recovered as pushing back the system to the “good” steady state requires that the parameter be decreased until the point F_2 , where the “bad” steady state disappears.

As a simple example from physics one may think of a glass of water, exhibiting a critical transition from liquid to solid when the temperature (which may be viewed as a “policy parameter”) varies and falls below 0° . In economics, a complex system consists of many economic agents (consumers, firms, investors, banks, etc.), which may be heterogeneous in various aspects. One would like to understand the emergent properties of complex macroeconomic systems and, in particular, how policy parameters might affect these emergent aggregate outcomes.

Perhaps the most crucial difference from complex systems in the natural sciences is that in economics and the social sciences, the “particles can think” and one needs a theory of adaptive behavior and learning.

In social-economic systems a theory of individual adaptive behavior is part of the law of motion of the macroeconomy. A central question to this survey is: what are the emergent properties of stylized complex macroeconomic systems with boundedly rational heterogeneous agents? Will a collection of boundedly rational heterogeneous agents be more likely to coordinate on the (homogeneous) rational outcome, or are fluctuation with booms and bust cycles a more likely aggregate outcome? This brings us back to Lucas (1986) (see the earlier quote) who views the question of collective behavior and coordination as an empirical question, an experimentally testable hypothesis. Indeed a large literature on *laboratory macro experiments* has developed in recent years, in

particular the learning-to-forecast experiments to study coordination of expectations in the lab. These macro experiments provide laboratory data, both at the individual (micro) and the aggregate (macro) level, which can be used to test, falsify, calibrate, or even estimate behavioral models. In this way, behavioral theory needs laboratory testing as a complementary tool for empirical analysis of various behavioral assumptions and models.

The survey is organized as follows. Section 2 discusses behavioral models with different degrees of (ir)rationality. There are many different models with boundedly rational interacting agents. To address the “wilderness of bounded rationality” our focus is on *parsimonious* decision rules that are validated in empirical work and laboratory experiments. This leads to stylized behavioral complexity models that are still partly analytically tractable.⁴ Section 3 discusses experimental macroeconomics and policy experiments, while section 4 summarizes and discusses policy implications of the observed coordination failure on nonrational, almost self-fulfilling equilibria.

2. Behavioral Models

What exactly is meant by “behavioral macroeconomics” is not easy to define. In his Nobel Prize lecture “Behavioral Macroeconomics and Macroeconomic Behavior,” Akerlof (2002) uses a very broad definition that, for example, includes models of asymmetric information, maintaining the assumption of

⁴Complementary to these stylized models, there is a large and rapidly increasing literature on agent-based simulation models using more detailed “bottom-up” modeling of individual decision rules of heterogeneous agents. The recent *Handbook of Computational Economics on heterogeneous agent modeling* (Hommes and LeBaron 2018) provides a state-of-the-art overview; see especially the survey by Dawid and Delli Gatti (2018) on agent-based macroeconomics.

rational expectations, to explain market failures.⁵ In a recent survey Driscoll and Holden (2014) summarize and discuss several concepts that behavioral economics has brought to macro-models, such as fairness considerations and other-regarding social preferences, cognitive biases, hyperbolic discounting of consumption and savings, habit formation, and rule-of-thumb consumption. De Grauwe (2012), in his *Lectures on Behavioral Macroeconomics*, emphasizes boundedly rational heterogeneous expectations in the new Keynesian macro-model, where agents switch between simple forecasting heuristics based upon their relative performance, as in Brock and Hommes (1997). There are thus many possible deviations—large or small—from the benchmark rational model. In the traditional macroeconomic paradigm there are (at least) three crucial assumptions underlying many models: (i) agents have rational expectations; (ii) agents behave optimally, that is, maximize utility, profits, etc.; and, related to both, (iii) agents have an infinite horizon for optimization and expectations. A pragmatic (but still admittedly subjective) definition of behavioral macroeconomics would be that (at least) one of these assumptions is relaxed and replaced by some form of bounded rationality. How many of these assumptions should be relaxed and by how much is then a matter of debate. For example, most agent-based models deviate from all of these three assumptions, to build a completely new macroeconomic system from “bottom-up” modeling of agents’ using simple micro-decision rules (heuristics); see Dawid and Delli Gatti (2018) for a recent survey on agent-based models in macroeconomics.

⁵Other approaches emphasizing informational frictions, but maintaining rational expectations, include rational inattention (Sims 2010) and imperfect knowledge (Angeletos and Lian 2016).

In our survey we focus on stylized behavioral models with learning and heterogeneous expectations. The question of what kind of (near) equilibria a population of heterogeneous boundedly rational forecasters might coordinate will play a prominent role throughout the survey. We start the survey with homogeneous adaptive learning (subsection 2.1), then move to heterogeneous expectations (subsection 2.2) and behavioral new Keynesian models (subsection 2.3).

2.1 Adaptive Learning

In the last three decades the adaptive learning approach has become a standard model of bounded rationality in macroeconomics. Agents behave as econometricians or statisticians and use an econometric forecasting model—the perceived law of motion—whose parameters are updated over time, for example, through recursive ordinary least squares, as additional observations become available. Early papers in this area are, for example, by Marcet and Sargent (1989a, b). The comprehensive overviews given by Evans and Honkapohja (2001) and more recently by Evans and Honkapohja (2013) have contributed much to its popularity in macroeconomics; see also Sargent (1993) for an early stimulating discussion of bounded rationality and learning.

Early work stressed learning of the parameters of a correctly specified model, that is, a perceived law of motion of exactly the same form as the (simplest) rational solution, with agents learning the parameters over time. Such an analysis then provides a stability theory of rational expectations equilibria and an equilibrium selection device to determine which rational equilibria are stable. Stability under adaptive learning should be seen as a minimum requirement of a rational expectations equilibrium (REE), because without stability under learning, coordination of a population of adaptive agents on a rational equilibrium seems highly unlikely.

2.1.1 Stability under Learning

For readers not familiar with adaptive learning it is useful to discuss stability under learning in a basic example. Consider a simple linear law of motion of the economy with an endogenous state variable x_t driven by exogenous stochastic shocks y_t :

$$(1) \quad x_t = a + bx_{t+1}^e + cy_{t-1} + u_t,$$

$$(2) \quad y_t = d + \rho y_{t-1} + \varepsilon_t.$$

To be concrete, one may think of x_t as an asset price, whose evolution is affected by price expectations x_{t+1}^e and by an exogenous AR(1) dividend process y_t with autocorrelation parameter ρ , $0 < \rho < 1$. The simplest rational solution, called the *minimum state variable (MSV) solution*, is of the form

$$(3) \quad x_t = \alpha + \gamma y_{t-1} + u_t,$$

with the price given as a linear function of the exogenous fundamental shocks (dividends). Assume for the moment that the parameters α and γ are fixed. Given that all agents believe that x_t follows the *perceived law of motion (PLM)* (3) the *implied actual law of motion (ALM)* becomes

$$(4) \quad x_t = a + b\alpha + b\gamma d + (c + b\gamma\rho)y_{t-1} + u_t.$$

A rational expectations solution is then a fixed point of the mapping T , from the PLM (3) to the ALM (4), and must satisfy

$$(5) \quad T(\alpha, \gamma) = (a + b\alpha + b\gamma d, c + b\gamma\rho).$$

The fixed point of the T -map corresponds to an REE solution and is given by:

$$(6) \quad \alpha = \frac{a}{1-b} + \frac{bcd}{(1-b)(1-b\rho)},$$

$$\gamma = \frac{c}{1-b\rho}.$$

Adaptive learning means that agents learn the parameters α and γ of the PLM (3) using estimation techniques such as ordinary least squares, which may be written in a recursive form algorithm. A simple associated differential equation governs the stability of the adaptive learning process and is given by

$$(7) \quad \begin{cases} \frac{d\alpha}{d\tau} = T_1(\alpha, \gamma) - \alpha = a + \frac{bcd}{1-b\rho} + (b-1)\alpha \\ \frac{d\gamma}{d\tau} = T_2(\alpha, \gamma) - \gamma = c + (b\rho - 1)\gamma. \end{cases}$$

The REE in (3) is also a fixed point of this differential equation (7) and, in this example, it is a (locally) stable fixed point, when the parameters $b\rho < 1$. One of the main general results from the adaptive learning literature is the *E-stability principle* stating that an REE (i.e., a fixed point of the T -map) is locally stable under adaptive learning processes such as ordinary least squares (OLS), when it is a locally stable fixed point of the associated ordinary differential equation (ODE). In this particular example, when agents believe that the PLM is of the form (3) and learn the parameters through OLS, the learning process converges (locally) to the REE. E-stability should be viewed as a necessary condition for REE to be empirically relevant. If an REE is not E-stable, then coordination of a large population of adaptive agents on such an equilibrium seems highly unlikely.

But what happens if the agents believe in a different PLM than the MSV solution (3)? For example, to forecast the state variable x_t it seems natural to include its lagged value x_{t-1} . Assume that instead of (3), agents believe that the PLM is of the (slightly) more general form

$$(8) \quad x_t = \alpha + \beta x_{t-1} + \gamma y_{t-1} + u_t.$$

This is an example where the PLM is *overparameterized* with respect to the MSV rational solution. In a similar way one can extend

the T -mapping $T(\alpha, \beta, \gamma)$ and simple algebra yields for α and γ the same REE fixed point as in (6) together with $\beta = 0$.⁶ It can be shown that this REE fixed point is again E-stable. The adaptive learning process is therefore robust with respect to overparameterization of the PLM in (8) and the REE in (3) is called *strongly E-stable*. Another parsimonious and perhaps plausible possibility would be that agents believe that the PLM is of the simpler form

$$(9) \quad x_t = \alpha + \beta x_{t-1} + u_t,$$

that is, agents do not realize that x_t is driven by an exogenous fundamental process y_t , but simply forecast x_t by lagged observations x_{t-1} . This is a simple example of misspecification, where the PLM is different from the MSV solution. We will return to the important issue of misspecification in subsection 2.1.3.

2.1.2 Endogenous Fluctuations under Adaptive Learning

Early work stressed adaptive learning as an *equilibrium selection* device of REE and studied E-stability of rational equilibria in various models, for example, in an asset pricing model with informed and uninformed traders (Bray 1982), the cobweb model (Bray and Savin 1986), in a general class of linear stochastic models (Marcet and Sargent 1989b) and in linear models with private information (Marcet and Sargent 1989a).

Later work has shown that adaptive learning need not converge to a rational expectations equilibrium, but learning may induce endogenous (periodic or even chaotic) business cycle fluctuations. Examples include the *learning equilibria* in overlapping generations models (Bullard 1994; Grandmont

⁶There is an additional REE fixed point $\beta = 1/b$, representing rational bubble solutions; see Evans and Honkapohja (2001).

1985, 1998), learning to believe in chaos (Schönhofner 1999), the consistent expectations equilibria in nonlinear cobweb models (Hommes and Sorger 1998), the learning to believe in sunspots (Woodford 1990), and the exuberance equilibria (Bullard, Evans, and Honkapohja 2008).

Constant Gain Learning.—Adaptive learning typically generates slow learning of parameters, because standard recursive estimation algorithms give equal weight to all past observations. Consequently, the weight given to the most recent observation becomes smaller and converges to 0 as the number of observations goes to infinity. The vanishing weight given to the most recent observation typically has a stabilizing effect on the learning dynamics. An alternative parameter updating scheme is *constant gain learning*, giving a fixed weight (the gain coefficient) to the most recent observations. Constant gain learning is consistent with lab experiments and survey data, where subjects or forecasters typically give more weight to the most recent observations. Constant gain learning models often give a better fit to macro and financial data and are able to generate observed stylized facts in time series data, such as high persistence, excess volatility, and clustered volatility (Evans and Honkapohja 2001; Sargent 1993; Milani 2007, 2011; Branch and Evans 2010).

Bubbles and Crash Dynamics under Learning.—Branch and Evans (2011a) develop a simple linear mean-variance asset pricing model capable of generating bubbles and crashes when agents use constant-gain learning to forecast expected returns and the conditional variance of stock returns.⁷

⁷There is a large literature on periodically collapsing rational bubbles. Blanchard and Watson (1982) develop a theory of rational bubbles in which agents' (rational) expectations are influenced by extrinsic random variables

Agents can choose between a risk-free asset paying a fixed return r and a risky asset (say a stock) paying stochastic dividends. Denote y_t as the dividend payoff and p_t as the asset price. Agents are risk averse and assumed to be *myopic mean-variance maximizers*. The mean-variance demand z_{dt} is

$$(10) \quad z_{dt} = \frac{E_t^*(p_{t+1} + y_{t+1}) - (1 + r)p_t}{a\sigma_t^2}$$

where $E_t^*(p_{t+1} + y_{t+1})$ denotes the conditional expectation of $p_{t+1} + y_{t+1}$, a is the risk aversion, and σ_t^2 denotes agents' conditional expectations about the variance of excess returns $p_{t+1} + y_{t+1} - (1 + r)p_t$. The equilibrium price is derived from market clearing $z_{dt} = z_{st}$ and given by

$$(11) \quad p_t = \frac{1}{1 + r} [E_t^*(p_{t+1} + y_{t+1}) - a\sigma_t^2 z_{st}].$$

The term $a\sigma_t^2 z_{st}$ may be seen as a time-varying risk premium. Dividends y_t and the supply of shares z_{st} are assumed to follow simple independent and identically distributed (IID) stochastic processes. Assuming $\sigma_t^2 = \sigma^2$ at steady state, the rational fundamental price can be computed as the discounted sum of future dividends minus the time-varying risk premium, and is given by

$$p_t^* = \sum_{j=1}^{\infty} \beta^j E_t(y_{t+j}) - \beta \sum_{j=0}^{\infty} \beta^j a\sigma^2 E_t(z_{st+j}),$$

where $\beta = 1/(1 + r)$ is the discount factor. There is additionally a class of rational

whose properties are in line with historical bubble episodes. West (1987), Froot and Obstfeld (1991), and Evans (1991) construct rational bubbles that periodically explode and collapse. A controversial issue for rational bubbles is that the trigger for the bubble collapse is often modeled by an exogenous sunspot process. In the model of Branch and Evans (2011a) bubbles and crashes arise endogenously as self-fulfilling responses to fundamental shocks, arising from the adaptive learning of agents.

bubble solutions, which are given by adding to the fundamental solution a rational bubble term $\beta^{-t}\eta_t$, where η_t is an arbitrary martingale, i.e., $E_t\eta_{t+1} = \eta_t$. Since $0 < \beta < 1$ the rational bubbles are explosive. Branch and Evans (2011a) show that the fundamental solution is E-stable under learning, while the rational bubble solutions are unstable under learning.

In Branch and Evans (2011a) agents' perceived law of motion is of the simple linear AR(1) form

$$(12) \quad p_t = k + cp_{t-1} + \epsilon_t,$$

where ϵ_t is an IID noise term. This linear specification coincides with the general form of the rational bubble solutions. Adaptive learning then consists of a recursive ordinary least squares updating scheme for the two parameters k and c of the conditional mean forecast together with a recursive algorithm for the conditional variance σ_t^2 of excess returns. For both learning processes constant gains can be used. Recursive updating of both the conditional variance and the expected return implies several mechanisms through which learning impacts stock prices. Extended periods of excess volatility, bubbles, and crashes arise with a frequency that depends on the extent to which past data is discounted. A central role is played by changes over time in agents' estimates of risk. First, occasional shocks can lead agents to revise their estimates of risk in a dramatic fashion. A sudden decrease or increase in the estimated risk of stocks can propel the system away from the fundamental equilibrium and into a bubble or crash. Second, along an explosive bubble path, risk estimates tend to increase and can become high enough to lead asset demand to collapse and stock prices to crash. Third, under learning, estimates for stock returns will occasionally escape to random walk beliefs that can be viewed as a bubble regime in which stock prices exhibit

substantial excess volatility. In this regime, revisions of risk estimates play an important role in generating the movements of prices that sustain the random walk beliefs. In summary, risk in an adaptive learning with constant gain setting plays a key role in triggering asset price bubbles and crashes. These intuitive and plausible results provide insights into the mechanisms by which expectations, learning, and bounded rationality generate large swings in asset prices.

2.1.3 Misspecification Equilibria

Under adaptive learning the PLM will, in general, be misspecified, that is, the PLM is generally different from the ALM. This observation has led to the study of *misspecification equilibria* under learning (Evans and Honkapohja 2001, Sargent 1999, Branch and McGough 2005, see especially the stimulating survey in Branch 2006). The idea here is that the representative agent uses a simple, parsimonious PLM to learn about the unknown ALM of the economy. These simple learning equilibria may be a more plausible outcome of the learning process of a population of adaptive agents.

Different types of parsimonious misspecification equilibria have been proposed in the literature. An interesting class are the *natural expectations* (Fuster, Laibson, and Mendel 2010; Fuster et al. 2012; and Beshears et al. 2013), where agents use a simple parsimonious fixed (higher order) AR(p) rule in forecasting to explain the long-run persistence of economic shocks. Since the parameters are fixed, strictly speaking this does not fall under adaptive learning, but its parsimony makes natural expectations intuitive and plausible forecasting rules.

Branch (2006) considers adaptive learning where the PLM is *underparametrized*, because agents do not take all relevant exogenous shock processes into account in their PLM. These beliefs, however, satisfy a least squares orthogonality condition consistent

with John Muth's original hypothesis. The least squares orthogonality condition in these models imposes that beliefs generate forecast errors that are orthogonal to an agent's forecasting model; that is, there is no discernible correlation between these forecast errors and an agent's model. Under this interpretation, the orthogonality condition guarantees that agents perceive their beliefs as consistent with the real world. Thus, agents can have misspecified (i.e., not rational expectations (RE)) beliefs, but within the context of their forecasting model they are unable to detect their misspecification. An equilibrium between optimally misspecified beliefs and the stochastic process for the economy is called a restricted perceptions equilibrium (RPE).

Branch and Evans (2010) apply these ideas in a mean-variance asset pricing model, where both dividends and the supply of shares follow exogenous stochastic AR(1) processes. There are two types of agents, who have different types of misspecified underparametrized price forecasting models. One type has a price forecasting model only based on the AR(1) dividend process, while the other type forecasts prices only based on the AR(1) process for the supply of shares. The RPE requires that agents forecast in a statistically optimal manner. It is required that the forecast model parameters are optimal linear projections, that is, the belief parameters, satisfy least-squares orthogonality conditions. Within the context of their forecasting model, agents are unable to detect their misspecification. Of course, if they step out of their model and run specification tests, they could detect the misspecification. But real-time simulations show that the misspecification is hard to detect and, for finite time, agents may not be able to reject their underparameterized models. They then study a misspecification equilibrium with *intrinsic heterogeneity*, and fractions of the two types of agents based on

their relative performance, as in Brock and Hommes (1997). The model exhibits multiple misspecification equilibria (ME) and the real-time learning dynamics switch between these different equilibria mimicking clustered volatility in asset returns.

Branch and Evans (2011b) use a similar approach in a new Keynesian macro model and study monetary policy under learning. There are two types of exogenous shocks to the economy: cost push shocks to the new Keynesian Phillips curve (NKPC) and supply shocks to the investment–savings (IS) curve, both following exogenous stochastic AR(1) processes. The RE MSV solution of the economy is a linear function of both shocks. There are two types of agents in the economy, one type using forecasts based only on the demand shocks and a second type using forecasts based only on the supply shocks. Branch and Evans (2011b) demonstrate that, even when monetary policy rules satisfy the Taylor principle by adjusting nominal interest rates more than one for one with inflation, there may exist equilibria with *intrinsic heterogeneity*, where the two types of agents coexist. Under certain conditions, there may exist multiple misspecification equilibria. These findings have important implications for business cycle dynamics and for the design of monetary policy. Branch and Evans (2011b) then study the role that policy plays in determining the number and nature of misspecification equilibria.

2.1.4 Behavioral Learning Equilibria

The most crucial aspect of adaptive learning is probably the choice of the PLM. For a large population of adaptive agents being able to coordinate their beliefs, the *parsimony* of the PLM seems crucial. Hommes and Zhu (2014) introduced a particularly simple form of misspecification called behavioral learning equilibrium. The idea here is that for each variable to be forecasted in the economy agents use a simple

(misspecified) univariate AR(1) forecasting rule. A *behavioral learning equilibrium* (BLE) arises when the sample average and the first-order autocorrelations of the AR(1) rule coincide with the observed realizations. Hence, along a BLE the parameters of the AR(1) rule are not free, but pinned down by two simple observable statistics, the sample average and the first-order sample autocorrelation.⁸ Agents thus use the optimal AR(1) forecasting heuristics. Such a simple, parsimonious learning equilibrium may be a more plausible outcome of the coordination process of individual expectations in large complex socioeconomic systems. The use of simple low-order autoregressive rules to forecast has also been documented in laboratory experiments with human subjects (e.g., Assenza et al. 2014).

Hommes and Zhu (2014) apply the BLE concept in the simplest class of models, where the actual law of motion of the economy is a one-dimensional linear stochastic process driven by exogenous AR(1) shocks.⁹ Two important applications of this framework are an asset pricing model driven by AR(1) dividends and an NKPC with inflation driven by an AR(1) process for marginal costs.

The NKPC with inflation driven by an exogenous AR(1) process y_t is given by (Woodford 2003)

$$(13) \quad \begin{cases} \pi_t = \delta\pi_{t+1}^e + \gamma y_t + u_t, \\ y_t = a + \rho y_{t-1} + \varepsilon_t, \end{cases}$$

⁸The idea behind BLE originates from the consistent expectations equilibria in Hommes and Sorger (1998), where the beliefs about sample average and all autocorrelations β^k , for all lags k , coincide with the realizations. Lansing (2009, 2010), applies the idea of (first-order) consistent expectations in a new Keynesian framework.

⁹Hommes et al. (2019) recently extended the BLE concept to higher dimensional linear stochastic models and estimated BLE in the Smets–Wouters DSGE model.

where π_t is the inflation at time t , π_{t+1}^e is the subjective expected inflation at date $t+1$, y_t is the output gap or real marginal cost, $\delta \in [0, 1)$ is the representative agent's subjective time discount factor, $\gamma > 0$ is related to the degree of price stickiness in the economy, and $\rho \in [0, 1)$ describes the persistence of the AR(1) driving process. Variables u_t and ε_t are IID stochastic disturbances with zero mean and finite absolute moments with variances σ_u^2 and σ_ε^2 , respectively.

Under RE inflation π_t is a linear function of the fundamental driving process y_t . The REE therefore has the same persistence and autocorrelations as the fundamental shocks. Assume instead that agents are boundedly rational and do not recognize or do not believe that inflation is driven by output gap or marginal costs, and therefore do not recognize that inflation should be a linear function of the exogenous shocks. Rather, agents believe that inflation follows a stochastic AR(1) process and simply forecast inflation by a (two-period ahead) univariate AR(1) rule, i.e., $\pi_{t+1}^e = \alpha + \beta^2(\pi_{t-1} - \alpha)$. The implied ALM then becomes

$$(14) \quad \begin{cases} \pi_t = \delta[\alpha + \beta^2(\pi_{t-1} - \alpha)] + \gamma y_t + u_t, \\ y_t = a + \rho y_{t-1} + \varepsilon_t. \end{cases}$$

Hommes and Zhu (2014) compute the corresponding first-order autocorrelation coefficient $F(\beta)$ of the implied ALM (14) as

$$(15) \quad F(\beta) = \delta\beta^2 + \frac{\gamma^2\rho(1 - \delta^2\beta^4)}{\gamma^2(\delta\beta^2\rho + 1) + (1 - \rho^2)(1 - \delta\beta^2\rho) \cdot \frac{\sigma_u^2}{\sigma_\varepsilon^2}}.$$

and show that there exists at least one non-zero BLE (α^*, β^*) with $\alpha^* = \bar{\pi}^*$ (i.e., the sample average equals REE inflation) and β^* a fixed point of the autocorrelation map $F(\beta)$ in (15).

Hommes and Zhu (2014) also show that when $F'(\beta^*) < 1$ the E-stability principle holds for the sample autocorrelation (SAC) learning process to learn the optimal parameters α^* and β^* . The time-varying parameters are given by the sample average

$$(16) \quad \alpha_t = \frac{1}{t+1} \sum_{i=0}^t x_i,$$

and the first-order SAC coefficient¹⁰

$$(17) \quad \beta_t = \frac{\sum_{i=0}^{t-1} (x_i - \alpha_t)(x_{i+1} - \alpha_t)}{\sum_{i=0}^t (x_i - \alpha_t)^2}.$$

Interestingly, for the NKPC multiple BLE may coexist, because the nonlinear autocorrelation map $F(\beta)$ may have multiple fixed points. Figure 2 illustrates the coexistence of a low- and a high-persistence BLE, which are both stable under SAC learning for appropriate initial states. The low-persistence regime represents a rather stable economy with inflation close to target, while the high-persistence regime is rather unstable with long-lasting periods of high or low inflation. The high-persistence BLE is characterized by $\beta^* \approx 0.996$, very close to unit root, and thus exhibits *persistence amplification*, with much more persistence in inflation than under RE. Under SAC learning with constant gain, the economy may switch irregularly between phases of low and high persistence and volatility in inflation.

This example shows how a very simple form of misspecification may lead to multiple equilibria and tipping points or critical transitions (compare figure 1 in the introduction to figure 2, panel H) between different regimes of low volatility and low persistence to high volatility and high

persistence. For initial states close to the target, SAC learning converges to the low persistence BLE. For initial states further away from the target, SAC learning converges to the high-persistence BLE. Under constant gain learning, the system may switch between both BLE. These results are consistent with the empirical finding in Adam (2007) that the restricted perception equilibrium (RPE) describes subjects' inflation expectations surprisingly well and provides a better explanation for the observed persistence of inflation than REE. Multiplicity of learning equilibria leaves an important task for monetary policy to keep inflation and output in the low-volatility regime. This simple model also shows how a simple and plausible form of misspecification brings us from a perfect rational world with a unique equilibrium into a more realistic complex boundedly rational reality with multiple equilibria and critical transitions.

2.1.5 Policy under Adaptive Learning

If coordination of a population of agents is better described by an adaptive learning process than by a rational expectations equilibrium, this has important policy implications. This subsection discusses some examples of policy analysis under models of adaptive learning.

In rational expectations models one can distinguish between *determinacy* and *indeterminacy* of equilibria. An REE is determinate when there exists a unique solution, typically a saddle-path solution converging to the rational steady state. An REE is indeterminate when multiple (typically a continuum) of solutions converging to the steady state exist. In such a case, often additional sunspot equilibria exist. If an REE is determinate, it is usually *assumed* that agents coordinate on the unique saddle-path solution. Such a saddle-path solution usually can only be computed by advanced computational software, such as the widely used

¹⁰An important and convenient feature of this natural learning process is that $-1 \leq \beta_t \leq 1$, since it is a (first-order) autocorrelation coefficient (Hommes and Sorger 1998).

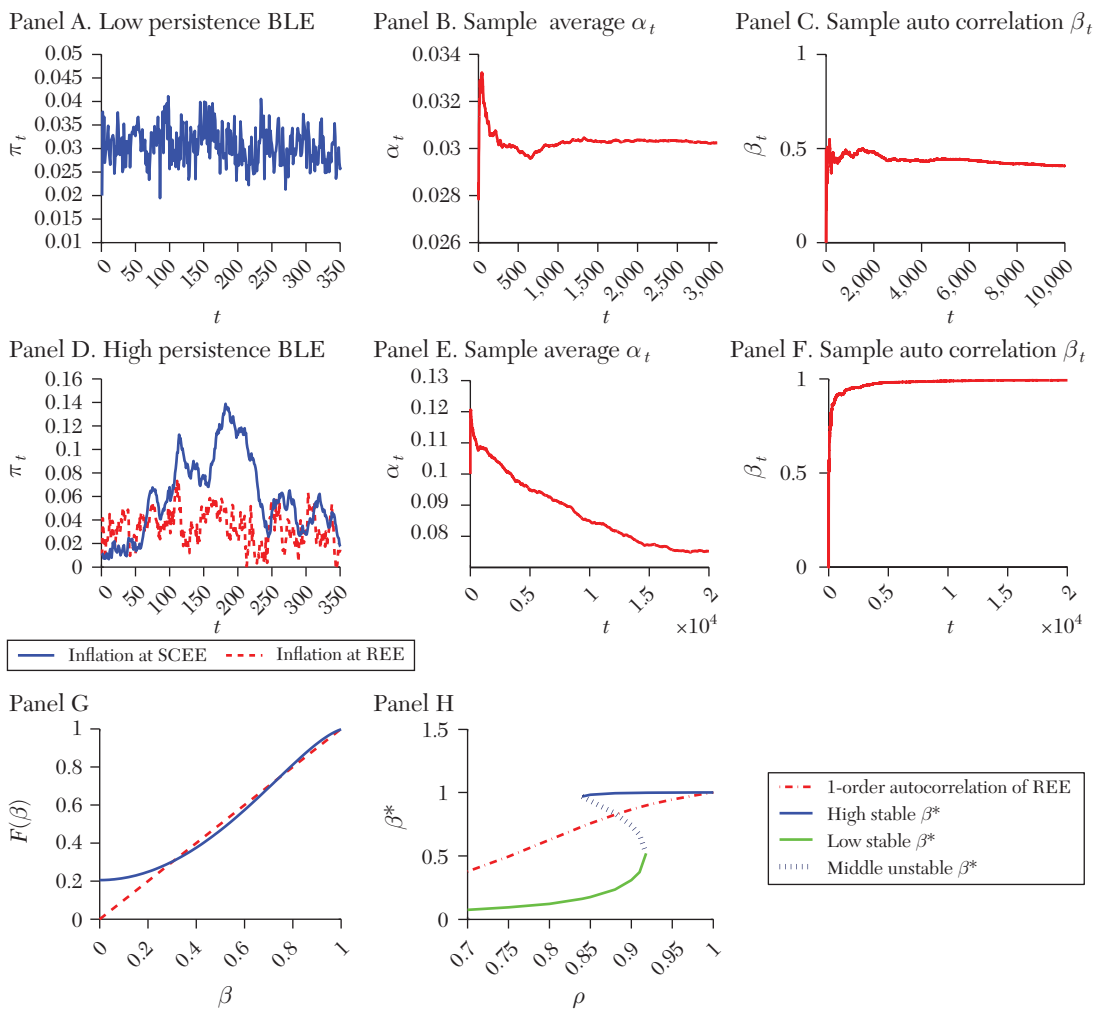


Figure 2. Multiple Behavioral Learning Equilibria in the NK Model

Notes: Top panels: convergence of SAC learning to low-persistence BLE $(\alpha^*, \beta_1^*) = (0.03, 0.3066)$. Middle panels: convergence to high-persistence BLE $(\alpha^*, \beta_3^*) = (0.03, 0.9961)$ exhibiting persistence amplification (for REE autocorrelation is $\rho = 0.9$). Panel G: BLE β^* correspond to the three fixed points of autocorrelation map $F(\beta)$ in (15). Panel H: BLE as a function of autocorrelation parameter ρ of the shocks; more persistent shocks lead to critical transition to persistence amplification (Hommes and Zhu 2014).

Dynare software, assuming that the equations of the economy are common knowledge. A learning theory of coordination on a saddle-path equilibrium, without the demanding assumption of perfect knowl-

edge of the law of motion of the economy, is however lacking, and without an adaptive learning process, coordination of a population of individuals on an equilibrium, even if it is unique, seems unlikely.

Bullard and Mitra (2002) study monetary policy under adaptive learning of the MSV solution in the new Keynesian model and show that considering learning generally can alter the evaluation of alternative policy rules. The (log-linearized) NK model is given by Clarida, Galí, and Gertler (1999) and Woodford (2003)

$$(18) \quad x_t = \tilde{E}_t x_{t+1} + \frac{1}{\sigma} (\tilde{E}_t \pi_{t+1} - i_t) + u_t,$$

$$(19) \quad \pi_t = \kappa x_t + \delta \tilde{E}_t \pi_{t+1} + v_t,$$

where x_t is the output gap, π_t inflation, i_t the nominal interest rate, $\tilde{E}_t x_{t+1}$, $\tilde{E}_t \pi_{t+1}$ are expectations about next period's output gap and inflation, and u_t and v_t are exogenous shocks following AR(1) processes. Equations (18) and (19) represent the IS curve and the Phillips curve. Here δ is the discount factor, and

$$(20) \quad \kappa = \frac{(\sigma + \eta)(1 - \omega)(1 - \delta\omega)}{\omega},$$

with σ and η the inverses of, respectively, the elasticity of intertemporal substitution and the elasticity of labor supply, and $(1 - \omega)$ is the fraction of firms that can adjust their price in a given period. Expectation \tilde{E}_t follows an adaptive learning process of the MSV solution $y_t = a + by_{t-1} + cw_t$, where $y_t = [x_t, \pi_t]^T$ and $w_t = [u_t, v_t]^T$.

The nominal interest rate is set by the central bank and Bullard and Mitra (2002) consider three different specifications of the Taylor interest rate rule, where the interest rate is set in response to deviation of inflation and output gap from the targets:

$$(21) \quad i_t = \phi_\pi \pi_t + \phi_x x_t \quad (\text{contemporaneous})$$

$$(22) \quad i_t = \phi_\pi \pi_{t-1} + \phi_x x_{t-1} \quad (\text{lagged})$$

$$(23) \quad i_t = \phi_\pi \tilde{E}_t \pi_{t+1} + \phi_x \tilde{E}_t x_{t+1} \quad (\text{forward looking})$$

where the coefficients $\phi_\pi, \phi_x > 0$ determine how strongly the central bank (CB) responds to inflation and output gap respectively.

Bullard and Mitra (2002) show that for the contemporaneous interest rate rule the determinacy (indeterminacy) region under RE coincides exactly with the E-stability (E-instability) region under learning. In this case, the policy analysis under RE and adaptive learning of the MSV solution are the same. For the forward-looking and the backward-looking Taylor rules, however, these regions do not coincide, and determinacy under RE does not imply E-stability under learning. This stresses the fact that policy should be based on plausible and empirically relevant models of adaptive learning. For all policy rules the Taylor principle holds under learning, that is, adjusting the nominal interest rates more than one-for-one in response to inflation above target implies learnability. In subsection 3.2 we will return to this issue and discuss some laboratory experiments to test the validity of the Taylor principle. Bullard and Mitra (2002) stress the general point that learnability should be a necessary additional criterion for evaluating alternative monetary policy rules.

Monetary and Fiscal Policy in a Non-linear NK Model.—Evans, Guse, and Honkapohja (2008) and Benhabib, Evans, and Honkapohja (2014) study the a non-linear NK model with a zero lower bound (ZLB) on the interest rate under adaptive learning of the steady state. In this nonlinear NK model, two steady states may coexist, the target steady state and a ZLB steady state, and liquidity traps or deflationary spirals may arise. The nonlinear equations describing aggregate dynamics are given by

$$(24) \quad c_t = c_{t+1}^e \left(\frac{\pi_{t+1}^e}{\beta R_t} \right)^{1/\sigma},$$

$$(25) \quad \pi_t(\pi_t - 1) = \beta \pi_{t+1}^e (\pi_{t+1}^e - 1) \\ + \frac{v}{\alpha \gamma} (c_t + g_t)^{\frac{1+\epsilon}{\alpha}} \\ + \frac{1-v}{\gamma} (c_t + g_t) c_t^{-\sigma},$$

$$(26) \quad R_t = \begin{cases} 1 + (R^* - 1) \left(\frac{\pi_{t+1}^e}{\pi^*} \right)^{\frac{\phi_\pi R^*}{R^* - 1}} \left(\frac{c_{t+1}^e}{c^*} \right)^{\frac{\phi_y R^*}{R^* - 1}} \\ \text{if } \pi_t \geq \tilde{\pi} \\ \tilde{R} \\ \text{if } \pi_t < \tilde{\pi}. \end{cases}$$

Equation (24) describes the dynamics of net output c_t (i.e., output minus government spending) through a standard Euler equation, where c_{t+1}^e and π_{t+1}^e denote respectively expectations of future net output and inflation, R_t is the nominal gross interest set by the central bank, $0 < \beta < 1$ is the discount factor, and $\sigma > 0$ refers to the intertemporal elasticity of substitution.

Equation (25) is an NKPC describing the dynamics of inflation π_t , where g_t is government spending of the aggregate good, $\epsilon > 0$ refers to the marginal disutility of labor, $0 < \alpha < 1$ is the return of labor in the production function, $\gamma > 0$ is the cost of deviating from the inflation target under Rotemberg price adjustment costs, and $v > 1$ is the elasticity of substitution between differentiated goods. The term $\pi_t(\pi_t - 1)$ in equation (25) arises from the quadratic form of the adjustment costs. Let $Q_t \equiv \pi_t(\pi_t - 1)$. The appropriate root for given Q is $\pi \geq 1/2$, so one needs to impose $Q \geq -1/4$ to have a meaningful definition of inflation.

Equation (26) describes an *aggressive* monetary policy, where $\tilde{R} = 1.0001$ corresponds to the ZLB on the nominal interest rate.¹¹ The forward-looking monetary policy

rule (26) is defined as aggressive since, while in “normal” times ($\pi_t \geq \tilde{\pi}$) it follows a standard forward-looking Taylor rule, it preventively cuts the nominal interest rate to the ZLB each time inflation drops below a given threshold $\tilde{\pi}$. The reaction coefficients in the interest rate rule are set to $\phi_\pi = 2$ and $\phi_y = 0.5$, which are in line with empirical estimates. This parametrization ensures local determinacy of the targeted steady state (π^*, c^*) under RE. However, as emphasised by Benhabib, Schmitt-Grohé, and Uribe (2002), “active” Taylor rules imply the existence of a second low-inflation steady state (π_L, c_L) , which is locally indeterminate under RE.

Fiscal policy is specified as

$$(27) \quad g_t = \bar{g},$$

where \bar{g} is fixed. Evans, Guse, and Honkapohja (2008) set $\pi^* = 1.05$ which implies a net output steady state value of $c^* = 0.7454$. Under the aggressive monetary policy in equation (26), the low-inflation steady state is given by $(\pi_L, c_L) = (0.99, 0.7428)$. The two equilibria of the model are depicted in figure 3. Evans, Guse, and Honkapohja (2008) consider a *fiscal switching rule* that can prevent liquidity traps and deflationary spirals. The fiscal switching rule prescribes an increase in public expenditures g_t each time monetary policy fails to achieve $\pi_t > \tilde{\pi}$. In model (24)–(25), given expectations π_{t+1}^e and c_{t+1}^e , any level of inflation π_t can be achieved by setting g_t sufficiently high. The idea behind the monetary–fiscal policy mix is the following. If the inflation target is not achieved under a standard Taylor rule, monetary policy is first relaxed in order to stimulate the economy. If the ZLB constraints the effectiveness of monetary policy, aggressive fiscal policy is then activated. As shown by Evans, Guse, and Honkapohja (2008), setting $\pi_L < \tilde{\pi} < \pi^*$ ensures the uniqueness of the targeted steady state. The unique equilibrium of the system

¹¹ $\tilde{R} > 1$ so that the corresponding interest rate $\tilde{R} - 1$ is small but positive at the ZLB.




Figure 3. Steady states in the NK model with ZLB.

Notes: Panel A: multiple equilibria with coexistence of low inflation steady state L and targeted steady state T under aggressive monetary policy. Panel B: unique equilibrium, that is, targeted steady state T under combined monetary policy and fiscal switching rule (Evans, Guse, and Honkapohja 2008).




Figure 4. Learning Dynamics under Aggressive Monetary Policy and Constant Fiscal Policy

Notes: Under adaptive learning the target steady state is locally stable, while the ZLB steady state is an unstable saddle point. Initial states below the stable manifold of the ZLB steady state fall into a liquidity trap or deflationary spiral under adaptive learning (Evans, Guse, and Honkapohja 2008).

under combined monetary (26) and fiscal policy (27) is illustrated in figure 3 (panel B).

E-stability and Equilibrium Selection.—The phase diagram of the dynamics under adaptive learning is given in figure 4. The solid black and the dashed black curves depict, respectively, the stable and unstable manifold of the low-inflation steady state saddle point (π_L, c_L) . The E-stability analysis shows that, although the targeted steady state is locally stable under learning, the saddle property of the low-inflation steady state creates a region in the phase space in which inflation and output decline over time. In particular, the stable manifold of the low inflation steady state divides the phase space in two regions: the *stable* region above the manifold, characterized by convergence to the targeted steady state (π^*, c^*) , and the *unstable* region below the manifold characterized by deflationary dynamics. This analysis shows that adverse expectational shocks may cause liquidity traps taking the form of deflationary spirals. Large pessimistic shocks may, in fact, push expectations into the unstable region, below the stable manifold of the low inflation steady state, leading to a self-reinforcing process in which inflation and output decline over time. On the other hand, when the aggressive monetary policy is augmented with the fiscal switching rule described in equation (27), the targeted steady state is globally stable under learning.

Evans, Guse, and Honkapohja (2008) thus show that an aggressive monetary policy rule alone can not recover the economy from a liquidity trap. Instead, an aggressive fiscal policy switching rule, in combination with an aggressive monetary policy that guarantees a lower bound on inflation, can recover the economy from a liquidity trap. Hommes, Massaro, and Salle (2019) design a lab experiment for empirically testing these predictions of the learning model in describing the occurrence of liquidity traps and monetary and

fiscal policies to recover the economy (see subsection 3.2).

Benhabib, Evans, and Honkapohja (2014) study fiscal and monetary policies in this NK model with a ZLB under infinite horizon learning. Unstable deflationary spirals occur after large expectational shocks. For large expectational shocks that push the interest rate rule to the ZLB a temporary fiscal stimulus or, in some cases a policy of fiscal austerity, can recover the economy from a deflationary trap when the policy is tailored in magnitude and duration. A fiscal stimulus “switching rule,” which automatically kicks in without discretionary fine-tuning, can be equally effective.

2.1.6 Internal Rationality

Adam and Marcet (2011) introduce a rather sophisticated notion of bounded rationality called *internal rationality*. Agents are *internally rational*, that is, they maximize discounted expected utility under uncertainty over an infinite horizon given dynamically consistent subjective beliefs about the future. But agents may not be *externally rational*, that is, they may not know the true stochastic process for payoff relevant variables beyond their control. Adam and Marcet (2011) focus on *near-rationality* in the sense that the subjective beliefs of agents are not exactly equal to the objective density of external variables, but will be close to the beliefs under RE by giving agents a prior distribution centered around the correct RE. They show then that even though this is potentially a small deviation from RE beliefs (when the variance of the prior is small), the outcomes of the learning model can be quite different.

Internal rationality may be viewed as a (sophisticated) microfoundation of adaptive learning. Adam and Marcet (2011) demonstrate how ordinary least squares—the most widely assumed learning rule in the adaptive-learning literature—arises as the optimal way to update conditional expectations from

a complete and dynamically consistent set of probability beliefs within a specific model. They also stress the point that microfoundations of the model are informative about which beliefs matter for the equilibrium outcomes in models of learning. In the general setting heterogeneous agents and market incompleteness are included to ensure that there is a distinction between the agent's own decision problem, which is perfectly known, and market behavior, which is assumed to be known only imperfectly.

Adam and Marcet (2011) focus on an asset pricing model with infinitely lived risk-neutral investors, heterogeneous agents, and incomplete markets. It is remarkable that under internal rationality the competitive equilibrium price of this infinite horizon asset pricing model is not the discounted sum of expected future dividends, but rather reduces to the myopic one-period-ahead asset pricing model where the asset price equals the discounted sum of next period's subjective belief of total stock payoff, that is, (cf. equation (23) in Adam and Marcet 2011)

$$(28) \quad P_t = \delta E_t^P(P_{t+1} + Y_{t+1}).$$

This gives rise to a self-referential model of learning about prices that leads to momentum and mean-reversion behavior in asset prices. It is remarkable that this micro-founded derivation of the asset pricing model with learning reduces to exactly the same pricing equation as in the boundedly rational mean-variance behavioral asset pricing models that were introduced earlier.¹²

Adam, Marcet, and Nicolini (2016) estimate an asset pricing model with internally rational agents using quarterly US

price-to-dividend (PD) ratios 1927:II to 2012:II. Their learning model replicates a number of asset pricing stylized facts such as high volatility in PD ratios, high persistence in PD ratios, excess volatility in returns, and excess return predictability at longer horizons (e.g., five years). They stress that agents are nearly rational. The perceived distribution of price behavior, although different from the true distribution, is nevertheless close to it and the discrepancies are hard to detect. Adam, Marcet, and Beutel (2017) show that the subjective beliefs are also consistent with survey data on expected capital gains.

Whether or not learning matters has important policy implications, as the desirability of policy responding to asset price fluctuations will depend to a large extent on whether asset price fluctuations are fundamentally justified.

2.1.7 Anticipated Utility Approach

Modeling bounded rationality in an infinite horizon dynamic optimization setting is challenging.¹³ In an infinite horizon dynamic optimization setting, a departure from rational expectations requires an assumption about whether agents take into account that beliefs are updated over time. To do so requires that agents anticipate their updating of parameters under adaptive learning or their switching between rules in a heterogeneous expectations setting over an infinite time horizon. Such strong cognitive capabilities seem highly unlikely in the real economy. The behavioral assumption that beliefs are coming from a completed learning process and are perceived to be fixed is called the anticipated utility approach and forms the benchmark of the adaptive

¹²For example, equation (28) coincides with the behavioral asset pricing model with heterogeneous beliefs (Brock and Hommes 1998, equation 2.7), discussed in subsection 2.2.2 of this survey, as well as with the myopic mean-variance asset pricing model with adaptive learning of Branch and Evans (2011a) discussed in subsection 2.1.2.

¹³A finite horizon approach is perfectly feasible in overlapping generations and life cycle models. Bullard and Duffy (2001), for example, study learning equilibria and excess volatility in a life cycle economy with capital accumulation.

learning literature. Branch and McGough (2018) survey behavioral implementations of the anticipated utility approach. A first model is *shadow price learning* (Evans and McGough 2018), where agents do not know the value function and the transition dynamics, but instead make linear forecasts of the shadow price and the state. The shadow price approach provides behavioral rules consistent with two-period intertemporal optimization, but does not require the full sophistication needed for agents to solve the complete dynamic programming problem. Interestingly, the shadow price learning approach may sometimes converge to the fully optimal and rational solution, but may also lead to rich learning dynamics. An alternative approach is *Euler equation learning* (Honkapohja, Mitra, and Evans 2013), where agents make decisions based on their perceived Euler equation derived from intertemporal optimization, that is, agents make choices by equating expected marginal benefits with expected marginal costs.

The shadow price and the Euler equation approaches are based on one-step-ahead forecasts. Branch, Evans, and McGough (2013) develop an N -step-ahead Euler equation learning approach, where agents forecast their terminal asset position to solve their N -period consumption–savings problem.¹⁴ Letting the planning horizon N go to infinity (and imposing the transversality condition) leads to the infinite horizon learning approach, recently surveyed in Eusepi and Preston (2017). In the infinity horizon learning approach agents are optimizing anticipated utility maximizers. It is interesting to note that these bounded rationality approaches to adaptive learning ultimately yield very similar asset pricing and macro dynamics as the internal rationality approach discussed before.

¹⁴See also Woodford (2019) for a recent analysis of monetary policy when planning horizons are finite.

2.2 Heterogeneous Expectations

In the last two decades a large behavioral literature on heterogeneous agents models (HAMs) with boundedly rational agents having heterogeneous expectations has developed; see, for example, the extensive survey by Hommes (2006) and, more recently, the survey by Dieci and He (2018). The main learning mechanism here is a form of evolutionary selection among different forecasting models with agents gradually switching to better-performing rules (Brock and Hommes 1997).¹⁵ This approach has been inspired by more complex genetic algorithm (GA) simulation models, but focuses on more stylized, partly analytically tractable models. These switching models generate endogenous boom and bust cycles mimicking stylized facts from real macro-financial data, such as bubbles and crashes, high persistence, clustered volatility, and fat tails.

The HAM literature has also been inspired by the noise trader literature in finance, pioneered by DeLong et al. (1990a,b), who introduced models where one type of agents has rational expectations, while another type, the noise traders, have nonrational expectations. In the model of DeLong et al. (1990b) noise traders have a *misperception* of their expectation about next period's price of a risky asset. They show that noise traders can survive in the market and earn a higher expected return than rational traders. DeLong et al. (1990a) consider a noise trader model with positive feedback traders and show that, in the presence of positive feedback traders, rational speculation can be destabilizing. These examples go against the Friedman hypothesis that nonrational traders will be driven out of the market because

¹⁵There is also a related literature on learning in dynamic games, for example, reinforcement learning (Erev and Roth 1998) and experiences weighted attraction (EWA) learning (Camerer and Ho 1999).

they lose money against rational traders. Instead, these examples show that in a heterogeneous world nonrational traders can survive competition with rational agents.

A lucid early critique on the representative agent approach in macroeconomics is given in Kirman’s paper “Whom or What Does the Representative Individual Represent?” As an alternative, Kirman (1992) stressed the importance of agents’ *interactions* for the emerging aggregate behavior. Kirman (1993) proposed a stochastic model of recruitment through local interactions, based on Föllmer (1974) and more recently extended in Föllmer, Horst, and Kirman (2005). This “ant-model” is motivated from biology, describing how local interactions of ants may lead to an asymmetric distribution over two identical food sources. Kirman (1991) applied the “ant-model” to financial markets, where investors are either optimistic or pessimistic and form their opinions through local interactions. This leads to herding and bubble and crash dynamics in a financial market model with fundamentalists and chartists.

2.2.1 *Costly Rational versus Free Rule-of-Thumb Expectations*

In an influential paper, Brock and Hommes (1997) introduced a simple cobweb model with costly rational versus free naive expectations. A novel feature compared to the noise trader literature is an endogenous switching mechanism between strategies based upon their relative performance as measured by utility, profits, or forecasting performance. Agents switch between cheap but destabilizing naive expectations with prices moving away from steady-state equilibrium, and costly stabilizing rational expectations with prices converging back to (a neighborhood of) the steady-state equilibrium. This leads to highly irregular, chaotic price fluctuations, with the market switching back and forth between close to fundamental stable price

fluctuations and unstable price fluctuations with high volatility.

It is useful to discuss this example in some detail. Producers can either buy the rational expectations price forecast $p_{1,t}^e = p_t$, at positive information gathering costs C , or freely obtain the simple naive forecast $p_{2,t}^e = p_{t-1}$. In a cobweb economy with rational versus naive expectations, the market equilibrium price is determined by demand and aggregate supply of both groups, that is,

$$(29) \quad D(p_t) = n_{1t} S(p_t) + n_{2t} S(p_{t-1}),$$

where n_{1t} and n_{2t} represent the fractions of producers holding rational respectively naive expectations. Notice that rational agents have perfect foresight and therefore perfect knowledge about the market equilibrium equation (29). Hence, rational agents not only have exact knowledge about prices and their own beliefs, but in a heterogeneous world they must also have perfect knowledge about expectations or beliefs of *all other* agents. For linear demand and supply market clearing in this two-type cobweb economy gives

$$(30) \quad a - dp_t = n_{1t} cp_t + n_{2t} cp_{t-1}.$$

Fractions of rational and naive producers are endogenously updated over time according to evolutionary fitness. Agents tend to switch to strategies that have performed better in the recent past. The fractions of rational and naive producers are given by a discrete choice or multinomial logit model (Brock 1993, Blume and Easley 1993)

$$(31) \quad n_{ht} = \frac{e^{\beta U_{h,t-1}}}{Z_{t-1}},$$

where $Z_{t-1} = \sum e^{\beta U_{h,t-1}}$ is a normalization factor so that the fractions add up to 1 and β is the *intensity of choice* parameter, inversely related to the noise level, measuring how quickly agents switch to better performing

strategies. There are two extreme cases: (i) $\beta = 0$: random choice, all fractions have equal weight; and (ii) $\beta = \infty$: the “neoclassical limit,” *all* agents switch immediately to the best strategy according to realized fitness. The discrete choice model (31) thus reflects the idea that agents switch to better-performing strategies. The intensity of choice measures how fast agents switch to these better-performing strategies.

The discrete choice probabilities (31) are obtained from a random utility model (Anderson, Arrow, and Pines 1988; Manski and McFadden 1981) of the general form

$$(32) \quad U_{ht} = \pi_{ht} + S_{ht} + P_{ht} + \epsilon_{iht},$$

where π_{ht} is *private utility* (e.g., utility, profits, wealth, forecasting performance, etc.), S_{ht} is *social utility* (e.g., utility from social interactions, herding effects, mean opinion index, etc.), P_{ht} represents sensitivity to *policy variables* (e.g., policy shocks or policy announcements, etc.), and ϵ_{iht} is *idiosyncratic noise*. When the noise term has an extreme value distribution and the number of agents tends to infinity, the probability of selecting strategy h tends to the multinomial logit probabilities (31).

Brock and Hommes (1997) have stressed that the performance measure should consist of *observable* quantities. The general form of (32) includes social interaction effects as emphasized by Brock and Durlauf (2001a, b).¹⁶ In the general representation of utility (32), we have also included a term for policy effects to stress the potential to take forward-looking behavior into account of how strategies respond to (observable) policy announcements.

¹⁶Brock and Durlauf (2001a, b) have written extensive surveys on social interaction models in economics. A key feature of these models is that the social interaction effects lead to multiple steady states. Their approach leads to analytically tractable models that can be used in estimating social interaction effects in real data.

2.2.2 A Behavioral Asset Pricing Model

A large literature on behavioral fundamentalists’/chartists’ asset pricing models in finance has developed, pioneered by early contributions by Zeeman (1974) and Day and Huang (1990), empirically supported by survey data studies of Frankel and Froot (1986; 1987a, b). At the Santa Fe Institute an artificial agent-based stock market model has been developed in Arthur et al. (1997) and LeBaron, Arthur, and Palmer (1999). Brock and Hommes (1998) developed a simple, more tractable version of this type of model. In their behavioral asset pricing model with heterogeneous beliefs, agents switch between fundamentalists’ versus chartists’ strategies based upon relative profitability (cf. also Lux 1995). An important motivation for heterogeneous agents is the observed high trading volume in real markets, which is at odds with no-trade theorems in the standard representative rational agent models. HAMs mimic many stylized facts observed in real financial data, for example, excess volatility, boom–bust cycles, fat tails, clustered volatility, high trading volume correlated with volatility, etc. (see, e.g., the survey in LeBaron 2006).¹⁷

Here we discuss a stylized version of a behavioral asset pricing model with heterogeneous beliefs and the estimation of a two-type model on stock market S&P 500 data. A detailed derivation may be found in Brock and Hommes (1998) and Hommes (2013).

Investors can choose between a risk-free asset paying a fixed return r and a risky asset (say a stock) paying stochastic dividends. Denote y_t as the dividend payoff, p_t

¹⁷See also Adam et al. (2015) for a recent HAM with extrapolative beliefs that replicates the joint behavior of stock prices, trading volume, and investors’ expectations. They also study in how far a transaction tax may prevent stock price booms.

as the asset price, and $z_{h,t}$ as the number of shares bought by investor type h . Agents are assumed to be *myopic mean-variance maximizers*. Let E_{ht} and V_{ht} denote the “beliefs” or forecasts of trader type h about conditional expectation and conditional variance. The mean-variance demand z_{ht} of type h is

$$(33) \quad z_{ht} = \frac{E_{ht}[p_{t+1} + y_{t+1} - (1+r)p_t]}{aV_{ht}[p_{t+1} + y_{t+1} - (1+r)p_t]} \\ = \frac{E_{ht}[p_{t+1} + y_{t+1} - (1+r)p_t]}{a\sigma^2},$$

where a is the risk aversion and, for simplicity, the conditional variance $V_{ht} = \sigma^2$ is assumed to be equal and constant for all types. In the case of zero supply of outside shares the market clearing price is given by

$$(34) \quad p_t = \frac{1}{1+r} \sum_{h=1}^H n_{ht} E_{ht}[p_{t+1} + y_{t+1}],$$

where n_{ht} denotes the time-varying fraction of trader type h and $E_{ht}[p_{t+1} + y_{t+1}]$ denotes the beliefs about the future price and the future dividend by investor type h .¹⁸ Recall that the rational, fundamental price of the risky asset is the discounted sum of expected future dividends. Assume for the moment that the dividend process is IID, with mean \bar{y} , then the fundamental value $p^* = \bar{y}/r$ is constant. Brock and Hommes (1998) assume that all agents have correct beliefs about the *exogenous* dividend process, but heterogeneous beliefs about *endogenous* prices. Agents can compute the fundamental price based on dividends, but nevertheless believe, for example, because of strategic uncertainty

or idiosyncratic reasons, that price may deviate from its fundamental value. A convenient feature of these assumptions is that the model can be reformulated in terms of the deviation $x_t = p_t - p_t^*$ from the fundamental benchmark as

$$(35) \quad x_t = \frac{1}{R} \sum_{h=1}^H n_{ht} E_{ht} x_{t+1}.$$

The fractions of type h are given by the discrete choice model (31), as before, with the fitness measure equal to realized profits in period t , that is,

$$(36) \quad \pi_{ht} = (p_t + y_t - Rp_{t-1}) \frac{E_{h,t-1}[p_t + y_t - Rp_{t-1}]}{a\sigma^2} \\ = (x_t - Rx_{t-1} + \epsilon_t) \frac{E_{h,t-1}[x_t - Rx_{t-1}]}{a\sigma^2},$$

with $y_t = \bar{y} + \epsilon_t$ and ϵ_t is IID noise. Due to its simplicity the Brock–Hommes (1998) model can be handled (partly) analytically and it has been shown that the market becomes unstable when performance-based switching (driven by short run profitability) is sufficiently fast (i.e., when the intensity of choice is high). The behavioral model exhibits irregular bubble and crash dynamics with the market switching between unstable phases where trend-following strategies dominate and stable phases where fundamentalists dominate (see figure 5). This is a counterexample to the Friedman hypothesis, as fundamentalists, who believe in the rational fundamental price, are unable to drive out chartists from the market. Hence, when strategy switching behavior is driven by (short run) profitability, chartists are able to survive in the market. More generally, this type of simple heterogeneous agent asset pricing model is able to generate stylized facts of asset prices, such as bubbles and crashes, fat tails, and clustered volatility (Lux and Marchesi 1999, see also the survey in Lux 2009).

¹⁸Notice that in the homogeneous case, $H = 1$, and discount factor $\delta = 1/(1+r)$ this mean-variance asset pricing equation is exactly the same as the internal rationality pricing equation (28) of Adam and Marcet (2011) in subsection 2.1.6.

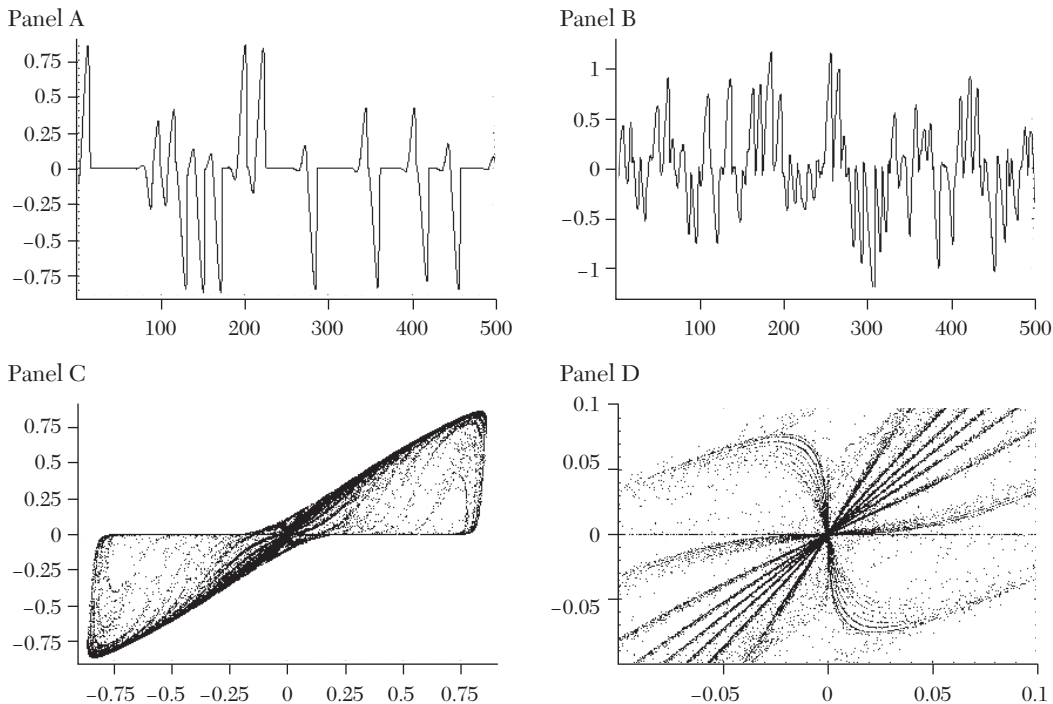


Figure 5. Bubble and Crash Dynamics in the Behavioral Asset Pricing Model with Heterogeneous Beliefs

Notes: Panel A: bubbles and crashes in deviations from the fundamental benchmark. Panel B: noisy bubble and crash dynamics. Panels C and D: strange attractors with chaotic dynamics in the phase space (x_t, x_{t-1}) . (Brock and Hommes 1998.)

More Hedging Instruments May Destabilize Markets.—Under the assumption of rational expectations futures markets or hedging instruments are usually stabilizing and welfare enhancing. Very little work has been done to study the role of futures and hedging instruments under bounded rationality. Brock, Hommes, and Wagener (2009) extended the behavioral asset pricing model to include hedging instruments in the form of Arrow securities. There are S states of the world arising with commonly known probabilities. There are n Arrow securities paying 1 if the corresponding state arises and 0 otherwise. For a sufficiently large number of Arrow securities, $n = S - 1$, the market is complete; for $n < S - 1$ the market is incomplete. Agents

are myopic mean-variance maximizers in this multi-asset world and have correct expectations on all stochastic dividends. Moreover, agents hold a common fundamental risk perception based on the variance-covariance matrix V_n of the dividend payoffs of the risky asset and the n Arrow securities. Brock, Hommes, and Wagener (2009) show that increasing the number of Arrow securities decreases the perceived fundamental risk and therefore boundedly rational agents take more leveraged positions, thus destabilizing the market, increasing price volatility, and decreasing average welfare. Figure 6 illustrates that adding more Arrow securities destabilizes the market and decreases average welfare. These destabilizing effects are relevant

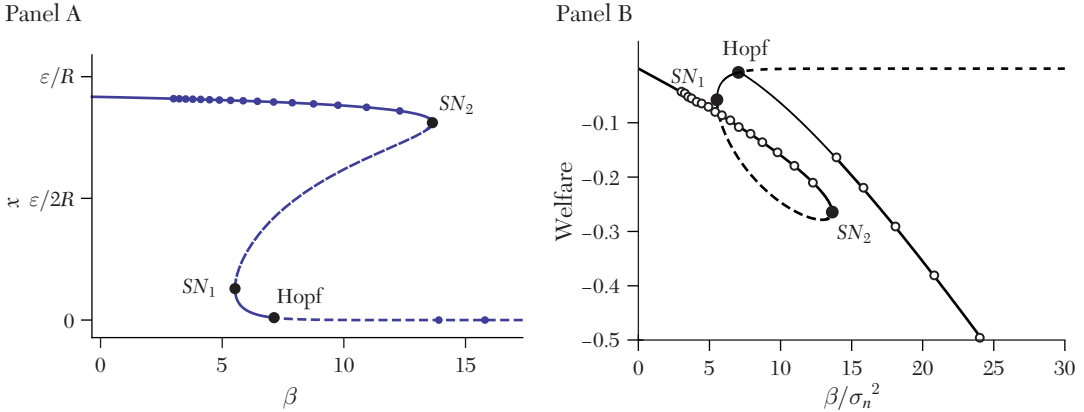


Figure 6. More Hedging Instruments May Destabilize Markets

Notes: Panel A: bifurcation diagram and critical transitions with respect to the intensity of choice β and the number of Arrow securities (the dots along the curves). As the number of Arrow securities increases, the system approaches the tipping point SN_2 along the non-fundamental steady state and then jumps to the fundamental steady state, which however already has become unstable after a Hopf bifurcation. In a boundedly rational world more Arrow securities thus destabilize the system. Panel B: average welfare decreases as a function of the intensity of choice parameter and the number of Arrow securities (Brock, Hommes and Wagener 2009).

to markets exposed to speculative trading and populated by boundedly rational agents.

2.2.3 Empirical Validation of Behavioral Asset Pricing Models

A large literature on empirical testing of behavioral heterogeneous agents models has developed, recently surveyed by Franke and Westerhoff (2017) and Lux and Zwickels (2018). Here we discuss the estimation of a two-type model in Hommes and in 't Veld (2017) and Boswijk, Hommes, and Manzan (2007); see also Lof (2015) and Chiarella, He, and Zwickels (2014) for similar models. In the behavioral asset pricing model, first a benchmark fundamental needs to be adopted. Empirically the dividend process is not IID, but the dividend data are well described by a geometric random walk with drift:

$$(37) \quad \log Y_{t+1} = \mu + \log Y_t + \nu_{t+1},$$

$$\nu_{t+1} \sim IID(0, \sigma_\nu^2).$$

The behavioral model then can be reformulated in price-to-cash flows. Investors have correct beliefs about dividends and estimate the constant growth rate $g \equiv e^{\mu + (1/2)\sigma_\nu^2}$ by averaging over $\log(Y_{t+1}/Y_t)$. Agents thus have model-consistent beliefs about the exogenous dividend process: $E_{i,t}[Y_{t+1}] = (1 + g)Y_t$. The RE fundamental price given by the discounted sum of expected future dividends, known as the Gordon model, is then given by

$$(38) \quad P_t^* = \frac{1 + g}{r - g} Y_t.$$

Hence, under RE the PD ratio is constant and given by

$$(39) \quad \frac{P_t^*}{Y_t} = \frac{1 + g}{r - g} \equiv \delta^*.$$

Figure 7 illustrates the S&P 500 stock market index, the fundamental value P_t^* , the PD ratio $\delta_t \equiv P_t/Y_t$ and the fundamental PD




Figure 7. Estimation of 2-Type Switching Model on S&P 500

Notes: Panel A: time series of S&P 500 and its fundamental value P_t^* . Panel B: PD ratio δ_t and its fundamental δ_t^* . Panel C: deviation of the PD ratio from its fundamental benchmark. Panel D: estimated fraction $n_{1,t}$ of fundamentalists. Panel E: the corresponding time-varying market sentiment ϕ_t in equation (48) (Hommes and in't Veld 2017).

ratio δ_t^* .¹⁹ The S&P 500 index clearly exhibits *excess volatility*, with much more volatility in asset prices than in the underlying fundamentals, a point already emphasized in the seminal paper of Shiller (1981).

Hommes and in 't Veld (2017) estimated the two-type model using quarterly data 1950: I–2016: IV. In deviations from the fundamental value $x_t \equiv \delta_t - \delta^*$, the two-type model can be rewritten as:

$$(40) \quad x_t = \frac{1}{R^*} (n_{1,t} E_{1,t}[x_{t+1}] + n_{2,t} E_{2,t}[x_{t+1}]),$$

$$R^* \equiv \frac{1+r}{1+g}.$$

The simplest form of heterogeneity occurs when belief types are linear in the last observation:

$$(41) \quad E_{h,t}[x_{t+1}] = \phi_h x_{t-1}.$$

These two types capture the only two possibilities in agents' beliefs: fundamentalists believe that the price will mean-revert back to its fundamental value ($0 \leq \phi_1 < 1$) and chartists believe that the price (in the short run) will move further away from the fundamental value ($\phi_2 > 1$).

The fractions of the two types are updated with a multinomial logit model as in Brock and Hommes (1997), with intensity of choice β :

$$(42) \quad n_{h,t+1} = \frac{e^{\beta U_{h,t}}}{\sum_{j=1}^H e^{\beta U_{j,t}}}.$$

¹⁹Boswijk, Hommes, and Manzan (2007) and Hommes and in 't Veld (2017) use the dynamic Gordon model with time variation in the interest rate and the growth rate of dividends. This dynamic approach is more flexible and allows for time variation in the fundamental PD ratio δ_t^* . As can be seen in figure 7 the time variation in the fundamental PD ratio of the dynamic Gordon model is relatively small. Notice that the dynamic Gordon model presupposes a fixed risk premium.

The performance measure $U_{h,t}$ is a weighted average of past profits $\pi_{h,t}$ and past fitness $U_{h,t-1}$, with memory parameter ω :

$$(43) \quad U_{h,t} = (1 - \omega)\pi_{h,t} + \omega U_{h,t-1},$$

with profits, up to a constant factor, given by

$$(44) \quad \begin{aligned} \pi_{h,t} &= z_{h,t-1} R_t \\ &= (\phi_h x_{t-2} - R^* x_{t-1})(x_t - R^* x_{t-1}), \end{aligned}$$

where $R^* = (1+r)/(1+g)$. The econometric form of the endogenous strategy switching model is an AR(1) model with a time-varying coefficient:

$$(45) \quad R^* x_t = n_{1,t} \phi_1 x_{t-1} + (1 - n_{1,t}) \phi_2 x_{t-1} + \epsilon_t,$$

$$R^* = \frac{1+r}{1+g},$$

where ϵ_t is an IID error term. Combining equations (42), (43), and (44), fractions depend nonlinearly on past realizations:

$$(46) \quad n_{1,t} = \left(1 + \exp \left[\beta (\phi_1 - \phi_2) \times \sum_{j=0}^{t-4} [\omega^j (1 - \omega) x_{t-3-j} \times (x_{t-1-j} - R^* x_{t-2-j})] \right] \right)^{-1},$$

$$(47) \quad n_{2,t} = 1 - n_{1,t}.$$

The estimated parameter values in Hommes and in 't Veld (2017) are:

- $\phi_1 = 0.948$: type 1 therefore are *fundamentalists*, expecting mean reversion of the price toward its fundamental value by 5.2 percent per quarter;
- $\phi_2 = 1.018$: type 2 are *trend extrapolators*, expecting the price deviation from

fundamental to increase by 1.8 percent per quarter;

- $\beta \approx 3.171$,²⁰
- $\omega = 0.824$: implying almost 20 percent weight is given to the most recent profit observation and about 80 percent to past profitability.

Define the *market sentiment* as

$$(48) \quad \phi_t = \frac{n_t \phi_1 + (1 - n_t) \phi_2}{R^*}.$$

Figure 7 (panels D and E) shows time series of estimated fractions of fundamentalists and the market sentiment. The fraction of fundamentalists varies considerably but gradually (due to memory) over time, with values between 25 percent and 90 percent until the 1990s, and more extreme values ranging from close to 0 to almost 100 percent after the dotcom bubble. The switching model offers an intuitive explanation of the dotcom bubble as being triggered by economic fundamentals (good news about a new internet technology) subsequently strongly amplified by trend-following behavior. Estimates of the market sentiment ϕ_t vary between 0.96 and 1 until the 1990s, showing near-unit root behavior. During the dotcom bubble the market sentiment ϕ_t exceeds 1 for several quarters and therefore the market is temporarily in an explosive bubble state. During the financial crisis the market is mainly dominated by fundamentalists indicating that the financial crisis has been reinforced by fundamentalists who expected a correction of asset prices back to fundamentals.

²⁰Estimating the β -parameter is hard, because of the highly nonlinear switching mechanism, and yields nonsignificant results due to the relatively small sample size. At the same time the coefficients ϕ_1 and ϕ_2 are significantly different from each other and therefore β is nonzero. See Hommes and 't Veld (2017) for bootstrap analyses and an extensive discussion.

In this behavioral asset pricing model with heterogeneous beliefs, agents switch between a mean-reversion and a trend-following strategy based upon realized profitability. Strategy switching driven by profitability leads to an almost self-fulfilling equilibrium with bubbles and crashes triggered by shocks (“news”) to economic fundamentals amplified by endogenous switching between trend-following and fundamentalists’ strategies.

Empirical Validation for Other Data Sets.—There is by now a large empirical literature estimating this type of heterogeneous agents model using various data sets, including stock prices, exchange rates, housing prices, and macro data (e.g. inflation); see, for example, Franke and Westerhoff (2017) and Lux and Zwickels (2018) for up-to-date surveys. A common finding is that bubbles are triggered by shocks to economic fundamentals (e.g., the dotcom bubble is triggered by a new internet technology) and strongly amplified by switching to (almost) self-fulfilling trend-following or chartists strategies. These trend-following strategies are profitable as long as the majority believes in them.

Heuristics switching models have also been applied to the housing market. Theoretical models for house price dynamics with heterogeneous expectations have been considered, for instance, by Dieci and Westerhoff (2012, 2013). Geanakoplos et al. (2012) develop an agent-based model to explain the housing boom and crash, 1997–2009, in the Washington DC area. Baptista et al. (2016) develop an agent-based model (ABM) of the UK housing market to study the impact of macro-prudential policies on key housing market indicators. Adam and Woodford (2012) consider a housing market model with Bayesian learning of an internally rational representative agent; Ascari et al. (2018) extend this model to the case of heterogeneous expectations with fundamentalists

versus chartists. Burnside, Eichenbaum, and Rebelo (2016) consider an epidemiological housing market model where agents disagree about the fundamental value of housing and infect each other.

Kouwenberg and Zwinkels (2014) estimated a HAM model specifically for the US housing market using quarterly data from 1960 until 2012. Ambrose, Eichholz, and Lindenthal (2013) examined a long time series of house price data of Amsterdam from 1650 to 2005 and found that substantial deviations from fundamentals persisted for decades and are corrected mainly through price adjustments and to a lesser extent through rent adjustments. Based on the same data set, Eichholtz, Huisman, and Zwinkels (2015) found that there is evidence for switching in expectation formation between fundamental and trend-following beliefs.

Bolt et al. (2019) estimate a two-type model—with fundamentalists versus trend-followers—to housing prices in deviations from a benchmark fundamental given by housing rents for eight different countries (United States, United Kingdom, the Netherlands, Japan, Spain, Switzerland, Sweden, and Belgium). Figure 8 illustrates some of the results for the United States, Japan, and the Netherlands. For all countries they find long-lasting and persistent explosive bubbles, with bubbles lasting sometimes more than a decade.

2.3 Behavioral New Keynesian Models

Behavioral asset pricing models originated more than two decades ago, as discussed in subsection 2.2.2. Behavioral macro models are of more recent date. For example, Gabaix (2017, 2020) recently introduced behavioral new Keynesian macro models, where full information infinite horizon optimization is replaced by *sparse dynamic optimization*, with agents ignoring or putting less weight on information in the distant future.

Gabaix (2017, 2020) mainly focuses on models where, given these informational restrictions, agents' behavior remains fully rational. This approach is similar in spirit to the rational inattention literature, surveyed in Sims (2010), where agents also partly ignore information or give less weight to some information but otherwise remain fully rational. The aggregate equations from Gabaix's (2020) behavioral NK model are:

$$(49) \quad x_t = ME_t x_{t+1} + \frac{1}{\sigma}(E_t \pi_{t+1} - i_t) + u_t, \quad (\text{IS curve})$$

$$(50) \quad \pi_t = \delta M^f E_t \pi_{t+1} + \kappa x_t + v_t, \quad (\text{Phillips curve})$$

$$(51) \quad i_t = \max\{\pi^T + \phi_\pi \pi_t + \phi_x x_t, 0\}, \quad (\text{Taylor rule})$$

where x_t is the output gap, π_t inflation, and i_t is the (contemporaneous) nominal interest rate rule. Compared to the NK benchmark model the crucial difference lies in the *attention parameters* $M, M^f \in [0, 1]$. The NK benchmark arises as a special case for $M = M^f = 1$. In the behavioral NK model agents are boundedly rational and are not fully forward looking, but less reactive to the future putting less weight on the far distant future. This is a form of cognitive discounting. Gabaix (2020) derives several policy implications within the behavioral model. For example, in the standard model, indeterminacy and multiplicity of equilibria arises when the Taylor principle does not hold. In contrast, in the behavioral model, whenever monetary policy is passive equilibrium is unique, even when the Taylor principle does not apply. The reason is that boundedly rational agents discount the future more and are thus less responsive to future events, lowering the complementarity between agents' actions. This force is stabilizing and

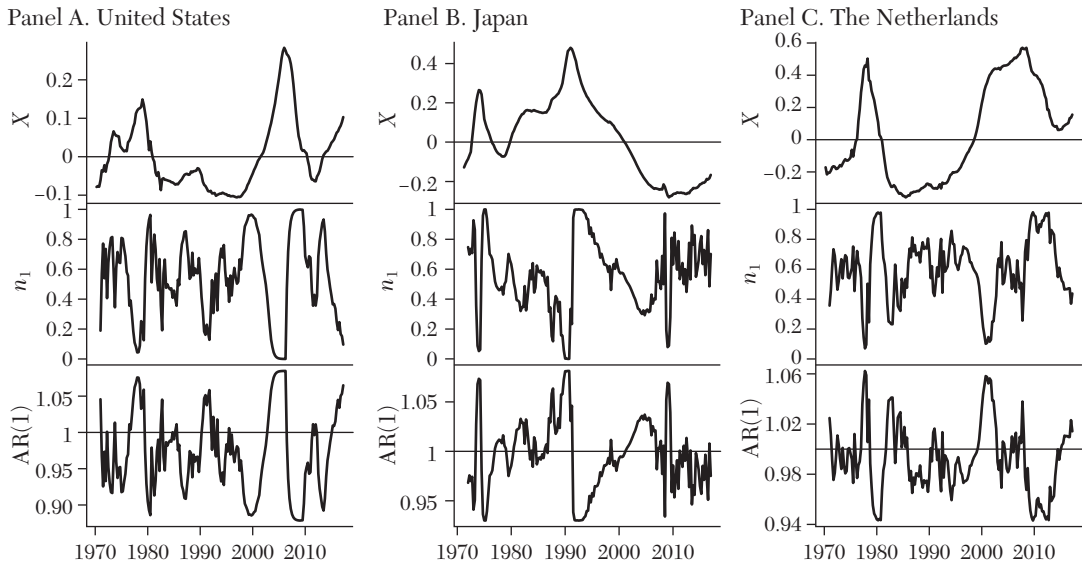


Figure 8. Estimation of 2-Type Switching Model on Housing Prices

Notes: Housing prices in the United States (left panels), Japan (middle panels), and the Netherlands (right panels). Top panels: relative house price deviations X_t from housing rent fundamentals; middle panels: estimated time-varying fractions of agents of type 1, that is, fundamental mean-reverting agents (middle panels); and bottom panels: estimated market sentiment as the time-varying AR(1) coefficient in equation (48). All countries show long and persistent temporary bubbles (Bolt et al. 2019).

dampens the possibility of multiple equilibria in Gabaix's (2020) behavioral NK model.

In this survey our focus is on the behavioral new Keynesian model with *nonrational heterogeneous expectations*. Even for linearized NK models, nonrational expectations and learning adds strong nonlinearities to the system giving rise for complex dynamics. It is interesting to note that, in contrast to Gabaix's behavioral model, nonrational expectations and learning introduce backward-looking expectations and typically add more *complementarities* and *positive feedback* to the NK framework making the dynamics of the model generally more unstable and allowing more easily for multiple equilibria and persistent deviations from the target steady state.

The heterogeneous expectations framework originates from Brock and Hommes (1997) with agents switching between different forecasting rules based upon their (recent) past relative performance. The learning mechanism here is thus characterized by "survival of the fittest," with agents gradually switching to better performing rules. Early applications of this heterogeneous expectations framework include Branch and McGough (2009, 2010), De Grauwe (2011, 2012), and Anufriev et al. (2013).

2.3.1 Heterogeneous Expectations

Branch and McGough (2009, 2010) are among the first papers to study new

Keynesian models with heterogeneous expectations, applying the framework of Brock and Hommes (1997) with agents switching between different forecasting rules based upon their (recent) past relative performance. Branch and McGough (2009, 2010) provide a micro-foundation for boundedly rational heterogeneous expectations at the agent level in the NK framework. They show that under Euler equation learning and a number of axioms about the individual forecasting rules, the aggregate IS and NKPC curves have the same functional form as under rational expectations. They consider an NK model with costly rational (perfect foresight) versus free naive expectations and show that instability and complicated dynamics may arise even if the model under rational expectations is determinate with a unique equilibrium path. When agents are allowed to switch between rules based upon their relative performance, complex dynamics (cycles and chaos) may arise even if the Taylor principle holds. Anufriev et al. (2013) consider a simple frictionless NK model with fundamentalists, optimists, and pessimists and show that multiple stable steady states coexist. They also consider a model with infinitely many different types and study the dynamics with a large type limit approximation (Brock, Hommes, and Wagener 2005). A 12-type model is shown to closely mimic US inflation time series, with large and highly persistent departures from target inflation. They also show that a more aggressive monetary policy Taylor rule reduces the number of steady states and can stabilize inflation.

Other early papers in this area are by De Grauwe (2011, 2012), who develops a behavioral NK macroeconomic model in which agents have cognitive limitations. Agents use simple but biased rules (heuristics) to forecast future output and inflation. Although the rules are biased, agents learn from their mistakes in an adaptive way, switching to better-performing strategies, as in Brock and Hommes

(1997). The model produces endogenous waves of optimism and pessimism (“animal spirits”) that are generated by the correlation of biased beliefs and match the stylized facts of inflation and output, such as persistence and fat tails. De Grauwe contrasts the dynamics of this model with a stylized DSGE version of the model and studies the implications for monetary policies. Strict inflation targeting is suboptimal, because it gives more scope for waves of optimism and pessimism to emerge, thereby destabilizing output and inflation.

Hommes and Lustenhouwer (2019b) analyze this NK model with optimists and pessimists in detail and also study the role of a ZLB on the interest rate. They show that multiple steady states exist, including a self-reinforcing liquidity trap steady state where the pessimistic agents dominate. They also consider a model with infinitely many types, using the concept of large type limit and show that self-fulfilling waves of optimism and pessimism may occur. More aggressive monetary policy and/or a higher inflation target reduce the number of steady states and make self-fulfilling animal spirits less likely.

The micro-foundations of heterogeneous expectations in the NK model have been discussed and studied in several papers. Branch and McGough (2009, 2010) used a one-step-ahead Euler equation approach to derive the aggregate IS and NKPC under heterogeneous expectations, after imposing restrictive assumptions about the individual forecasting rules. Massaro (2013) uses an infinite horizon optimization framework and derives aggregate IS and NKPC under heterogeneous expectations, with the homogeneous rational agent benchmark nested as a special case. Mauersberger (forthcoming) develops a micro-founded NK model where agents form forecasts of household relevant variables, as opposed to economy-wide aggregates, and test this model in the laboratory.

Kurz, Piccillo, and Wu (2013) explore an NK model with diverse beliefs and show that the aggregate IS and NKPC depend on an aggregate state variable named “mean market state of belief.” Diverse beliefs alter the problem faced by a central bank since the source of fluctuations is not only exogenous shocks but also market expectations. They show that due to diverse beliefs the effects of policy instruments are not monotonic and the trade-off between inflation and output volatility is complex. Monetary policy can counter the effects of market belief by aggressive anti-inflation policy, but at the cost of increased volatility of financial markets and individual consumption.

Hommes and Lustenhouwer (2019a) simplify the aggregate equations of Kurz, Piccillo, and Wu (2013) using a property of the discrete choice model for strategy switching of Brock and Hommes (1997). Under this model it is implicitly assumed that the probability to follow a particular heuristic next period is the same across agents, that is, independent of the heuristic they followed in the past. This reflects the fact that agents are not inherently different, but face the same trade-off between heterogeneous forecasting rules. They assume agents know (have learned) that all agents have the same probability to follow a particular heuristic in the future, and that they know that consumption decisions only differ between households in so far as their expectations are different. In that case households expectations about their own future consumption coincide with their expectations about the future consumption of any other agent. Agents therefore realize they should base their current period consumption decision on expectations about future aggregate consumption and, therefore, Kurz’s mean market state of belief reduces to 0. Under these assumptions aggregate equations of the NK model with switching among heterogeneous expectations rules are equivalent to

replacing conditional expectations by average expectations.

2.3.2 *Individual and Social Learning*

Another, closely related approach to learning with heterogeneous forecasting rules is called *individual* or *social learning* and uses genetic algorithm learning based on evolutionary selection, mutation, and crossover. The evolutionary approaches may be viewed as a more descriptive form of actual learning behavior in complex market economies. The notion of *individual evolutionary learning* (IEL) was introduced in Arifovic and Ledyard (2011), building on the work of Arifovic (1994), as a way of modeling heterogeneous strategies in large strategy spaces with a continuum of decision choices (as opposed to reinforcement learning (Erev and Roth 1998) and EWA learning (Camerer and Ho 1999) in game theoretic settings, where the strategy space is finite). Individual learning refers to agents evolving their own set of successful strategies, as opposed to a social learning process where the population of strategies evolves and each agent is represented by a single strategy.

Arifovic et al. (2010) apply IEL in an agent-based dynamic extension of the Kydland and Prescott (1977) model and study the role of cheap talk announcements by the policy maker. Private agents can choose between two strategies: *believe*, that is, act as if the policy announcement was true; or *not believe* and compute the best possible forecast of the policy maker’s next action. In each period word-of-mouth information exchange allows a fraction of agents to compare their last-period payoffs with the ones obtained by agents who followed the other strategy. Each agent then adopts the strategy that provided the highest payoff and uses it until a new comparison motivates it to switch strategies again. The proportion of believers thus may change over time and can be interpreted as a measure of the policy maker’s credibility.

Simulations show that the policy maker is able to learn how to reach an outcome that is Pareto-superior compared to the one that would be attained without adequate cheap talk. This outcome is characterized by a succession of trust building phases, where the announcement and the true inflation are chosen in order to increase the proportion of believers; and trust exploitation phases, where the policy maker uses the existence of a large fraction of believers to achieve for itself high payoffs at the cost of a decrease of the fraction of believers.

Arifovic, Bullard, and Kostyshyna (2013) analyze the effects of social learning in a new Keynesian monetary policy context. Social learning may be viewed as a more descriptive and realistic form of actual learning behavior in complex market economies. In this NK framework the Taylor principle governs uniqueness and expectational stability of REE under adaptive learning. Surprisingly, they find that the Taylor Principle is not necessary for convergence to REE MSV equilibrium under social learning. Under social learning for all policy parameters the system seems to converge to the REE. Sunspot equilibria also exist in the indeterminate region. Under social learning agents cannot coordinate on a sunspot equilibrium in general form specification, however, they can coordinate on common factor specifications. It remains unclear how general these stability results under social learning are and whether they are, for example, robust with respect to some of the underlying assumptions, such as the correct specification of the perceived law of motion as the MSV (no misspecification), the mutation rate is assumed to decrease over time (favoring the REE forecast) and there is no ZLB.

Hommes et al. (2017) study a GA model in the new Keynesian framework with an inflation targeting interest rate rule (with a ZLB), where the GA optimizes a simple first-order

forecasting heuristic as in Heemeijer et al. (2009). This means that the PLM is misspecified and in particular takes a trend-following coefficient into account. The stability of GA learning coincides with the Taylor principle; when the Taylor principle is not satisfied, the NK model is unstable under GA learning and may yield explosive inflationary or deflationary spirals and persistent fluctuations in inflation and output. When the Taylor principle holds GA learning is more stable, although some endogenous fluctuations and oscillatory behavior may arise. The GA simulations fit the different types of observed behavior—monotonic convergence, oscillatory behavior, and deflationary spirals (when the Taylor principle does not hold) in the NK laboratory experiments of Assenza et al. (2021) (see subsection 3.2).

Arifovic, Schmitt-Grohé, and Uribe (2018) study *social learning* in the NK model with a ZLB. There are three REE: the normal target steady state, the always binding ZLB steady state, and an occasionally binding ZLB steady state following a Markov process. Agents' PLM is a correctly specified rule capturing the three REE steady states. Mutation and crossover probabilities are assumed to be constant (10 percent) over time. It is also assumed that the current realizations of the stochastic process for the shocks to the natural rate can be observed. The ZLB steady state is unstable under adaptive learning, but interestingly it is stable under social learning, where agents learn the optimal coefficients of the PLM. Agents can learn to have pessimistic sentiments about the central bank's ability to generate price growth, giving rise to a stochastically stable environment characterized by deflation and stagnation.

2.3.3 *The Two-Type Switching Model and Forward Guidance*

Heterogeneous expectations provide a natural framework to study the role of the

ZLB and the credibility of the central bank to prevent or recover the economy from liquidity traps. Hommes and Lustenhouwer (2019a) introduce *endogenous credibility* of the central bank in an NK framework with heterogeneous expectations.

There are two types of agents, credibility believers and naive expectations. The first forecasting heuristic can be described as “trust the central bank.” Followers of this heuristic are called fundamentalists or credibility believers, and expect future inflation and output gap to be equal to the targets of the central bank. The fraction of fundamentalists can be interpreted as the *credibility* of the central bank. In contrast with rational expectations models, this model therefore involves *endogenous credibility*. The fundamentalists compete with naive expectations, which uses the last observation as a best guess for future realizations of inflation and output. Notice that the naive heuristic coincides with rational expectations when inflation or output follows a random walk. If inflation or output follows a near-unit root process, the naive forecast is therefore nearly rational. Naive agents furthermore add persistence in inflation and output gap to the model in a very simple and intuitive manner, without the need to assume heavily serially correlated shocks. There is empirical evidence matching the type of heterogeneity in the two-type model, both in survey data and lab experimental data, as will be further discussed in subsections 2.3.4 and 3.2.

The model is given by an NKPC describing inflation π_t , an IS curve describing output gap x_t , and a policy rule for the nominal interest rate i_t with a ZLB:

$$(52) \quad x_t = \tilde{E}_t x_{t+1} + \frac{1}{\sigma} (\tilde{E}_t \pi_{t+1} - i_t) + u_t,$$

$$(53) \quad \pi_t = \delta \tilde{E}_t \pi_{t+1} + \kappa x_t + v_t,$$

$$(54) \quad i_t = \max \left\{ \pi^T + \phi_1 (\tilde{E}_t \pi_{t+1} - \pi^T) + \phi_2 (\tilde{E}_t x_{t+1} - x^T), 0 \right\},$$

where \tilde{E}_t denotes aggregate expectations of all agents in the economy. Expectations are formed using the two simple heuristics, fundamentalists (type 1) and naive expectations (type 2), as

$$(55) \quad E_{1t} x_{t+1} = x^T \quad \text{and} \quad E_{1t} \pi_{t+1} = \pi^T, \\ \text{(fundamentalists)}$$

$$E_{2t} x_{t+1} = x_{t-1} \quad \text{and} \quad E_{2t} \pi_{t+1} = \pi_{t-1} \\ \text{(naive).}$$

Agents switch between these rules, and the fractions n_{1t} and n_{2t} of the two types are given by the discrete choice probabilities (31), with fitness based on the squared prediction errors. Aggregate expectations about inflation and output are then given by

$$(56) \quad \tilde{E}_t \pi_{t+1} = n_{1t}^{\pi} \pi^T + n_{2t}^{\pi} \pi_{t-1},$$

$$(57) \quad \tilde{E}_t x_{t+1} = n_{1t}^x x^T + n_{2t}^x x_{t-1}.$$

The fraction of fundamentalists is then the endogenous credibility of the CB and depends on how well the CB achieved its targets.

Hommes and Lustenhouwer (2019a) derive policy implications for an inflation targeting central bank, who's credibility is endogenous and depends on its past ability to achieve its targets. Interestingly, the region of allowed policy parameters is strictly larger under heterogeneous expectations than under rational expectations. Furthermore, with theoretically optimal monetary policy (with coefficients minimizing a quadratic loss function), global stability of the fundamental steady state can be achieved, implying that the system always converges to the targets of the central bank. This result, however, no longer holds when the ZLB on the nominal interest rate is accounted for. Self-fulfilling deflationary spirals can then occur, even

under optimal policy. The occurrence of these liquidity traps crucially depends on the credibility of the central bank. Deflationary spirals can be prevented with a high inflation target, aggressive monetary easing (i.e., cutting the interest rate to 0 when inflation falls below a threshold), or a more aggressive response to inflation (a higher coefficient in the Taylor rule). All these deviations from optimal policy have their costs, and may, for example, lead to higher output volatility, so a well balanced combination may be the best way to go.

Forward Guidance.—Goy, Hommes, and Mavromatis (2020) use the two-type switching model to study the macroeconomic effects of central bank forward guidance when central bank credibility is endogenous. In particular, they take a stylized new Keynesian model with an occasionally binding ZLB constraint on nominal interest rates and heterogeneous and boundedly rational households. The central bank uses a bivariate vector autoregression (VAR) model to forecast inflation and output gap. But their VAR model is misspecified, because it does not take into account the time variation in the distribution of aggregate expectations. In this framework, they introduce forward guidance by allowing the central bank to publish its own forecasts (Delphic guidance) and to commit to a future path of the nominal interest rate (Odyssean guidance). Both Delphic and Odyssean forward guidance increase the likelihood of recovery from a liquidity trap. While Odyssean guidance alone can increase ex post macroeconomic volatility and thus reduce welfare, it still appears to be more powerful.

To study forward guidance in a bounded rationality framework, Goy, Hommes, and Mavromatis (2020) use N -step-ahead Euler equation learning and heterogeneous expectations, so that future interest rate expectations matter. The NK model with

N -step-ahead Euler equation learning can be summarized as

$$(58) \quad x_t = \tilde{E}_t x_{t+N} - \frac{1}{\sigma} \tilde{E}_t \sum_{j=0}^{N-1} (i_{t+j} - \pi_{t+j+1} - \bar{r}) + u_t,$$

$$(59) \quad \pi_t = \delta^N \tilde{E}_t \pi_{t+N} + \tilde{E}_t \sum_{j=0}^{N-1} \delta^j \kappa x_{t+j} + v_t,$$

where \tilde{E}_t denotes the heterogeneous expectations operator to be specified below. The IS curve (58) and Phillips curve (59) pin down output x_t and inflation π_t , given a nominal interest rate i_t and N -step-ahead forecasts of the interest rate, inflation, and output gap. The term δ is the discount factor, \bar{r} is the steady state real interest rate, given by $\bar{r} = (1/\delta) - 1$, and u_t and v_t represent exogenous shocks. The model is closed using a contemporaneous Taylor-type rule with a ZLB:

$$(60) \quad i_t^{mp} = \max \left\{ 0, \bar{r} + \bar{\pi} + \phi \left(\pi_{t|t}^{e,cb} - \bar{\pi} \right) \right\},$$

where $\pi_{t|t}^{e,cb}$ denotes the central bank's real-time inflation projection, made at the beginning of period t . For the *unconventional* policy, we equip the central bank with two additional policy tools:

- Delphic guidance: to publish the CB forecasts $\pi_{t+j|t}^{e,cb}$, $i_{t+j|t}^{e,cb}$, and $x_{t+j|t}^{e,cb}$ for $j = 1, \dots, q^D$, or
- Odyssean guidance: commit to set interest rates $i_{t+j} = 0, \forall j = 1, \dots, q^O$.

There are two types of agents, credibility believers versus naive expectations. Let $z_t = x_t, \pi_t$ (i.e. output gap or inflation). Then expectations are given by:

$$(61) \quad \tilde{E}_{1,t} z_{t+j} = \begin{cases} z_{t+j|t}^{e,cb}, & \forall j = 1, \dots, q^D; \\ \bar{z}, & \forall j = q^D + 1, \dots, N; \end{cases}$$

and of naive expectations

$$(62) \quad \tilde{E}_{2,t} z_{t+j} = z_{t-1}, \quad \forall j = 1, \dots, N.$$

Notice that under forward guidance the credibility believers trust the central bank and use their forecast. The credibility of the central bank, however, depends on the fraction of credibility believers, which varies endogenously over time. Agents switch between credibility believers and naive expectation based on the relative forecasting performance, with fractions given by the discrete choice model (31).

It is assumed that households know the policy rule of the central bank and form expectations about the nominal interest rate using their own inflation expectations:

$$(63) \quad \tilde{E}_{1,t} i_{t+j} = \begin{cases} i_{t+j|t}^{e,cb} = \max\{0, \bar{r} + \bar{\pi} + \phi(\pi_{t+j|t}^{e,cb} - \bar{\pi})\}, \\ \quad \forall j = 1, \dots, q^D; \\ 0, \\ \quad \forall j = 1, \dots, q^O; \\ \bar{i}, \\ \quad \forall j = q^k + 1, \dots, N; \end{cases}$$

where Odyssean guidance is dominant in case of both, and

$$(64) \quad \tilde{E}_{2,t} i_{t+j} = \max\{0, \bar{r} + \bar{\pi} + \phi(\pi_{t-1} - \bar{\pi})\}, \\ \forall j = 1, \dots, N.$$

Finally, the central bank forms expectations by adaptive learning of a bivariate VAR(1) model

$$(PLM) \quad y_t^{e,cb} = A_0 + A_1 y_{t-1} \equiv A' w_{t-1}$$

where $y_t^{e,cb} \equiv [x_{t|t}^{e,cb}, \pi_{t|t}^{e,cb}]'$ and $w_{t-1} \equiv [1, x_{t-1}, \pi_{t-1}]'$. If the fractions of the two agent types were constant, the CB would have a correctly specified PLM. Since these fractions are time varying, the CB uses a misspecified model and we have an *RPE* (Branch 2006; see subsection 2.1.3).

Goy, Hommes, and Mavromatis (2020) show that the model has two steady states,

the target steady state, which is stable under learning, and the ZLB steady state, which is a saddle (see figure 9). Without forward guidance (FG) all initial states below the stable manifold of the ZLB steady state fall into a liquidity trap and exhibit a deflationary spiral. FG, however, enlarges the recovery region of the economy, as illustrated in figure 9. How much FG enlarges the recovery region of the economy depends critically on the credibility of the CB as measured by the fraction of credibility believers. When the fraction of credibility believers is small, FG is ineffective. On the other hand, when the fraction of credibility believers is large, FG has a large effect. This resolves the FG puzzle under rational expectations (Del Negro et al. 2012). In our model the effectiveness of FG depends critically on the credibility of the CB. Extensive Monte Carlo simulations show that without FG the probability of deflationary spirals is about 28.8 percent and with FG about 14.5 percent (Delphic FG), 10.8 percent (Odyssean FG), and 12.1 percent (both).

2.3.4 Empirical Validation of the Two-Type Switching Model

Both surveys of consumers and professional forecasters and laboratory experiments with human subjects show that there is considerable heterogeneity in inflation forecasts consistent with the two-type model discussed above.

The case study of the Volcker disinflation by Mankiw et al. (2004) nicely illustrates the presence of two types of heuristics in survey data. In figure 10 (Mankiw et al. 2003, p. 237) the evolution of inflation expectations as measured by the Michigan Survey from 1979 up to and including 1982 is plotted. They show that at the start of 1979 expectations were centered around a high inflation value. Over the next eight quarters (during which Paul Volcker was appointed chairman of the Board of Governors of the Federal Reserve

Panel A. Scenario I: All naive in t and $t + 1$

Panel B. Scenario II: 50% naive in t and $t + 1$

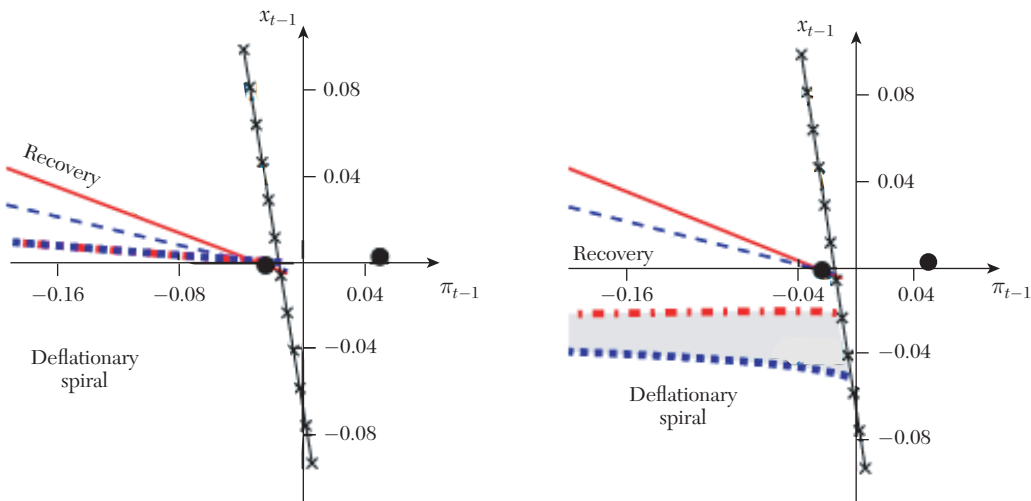


Figure 9. Recovery Regions under FG

Notes: The effectiveness of FG depends crucially on the credibility of the CB. The figure shows the basin of attraction of the target steady state (right black dot at $(\pi^*, x^*) = (\bar{\pi}, \bar{x})$). For scenario I (panel A), we assumed all households to be naive in periods t and $t + 1$, while for scenario II (panel B), the fractions are equal. The (left) black dot at $(\pi^*, x^*) = (-\bar{r}, -(1 - \delta)/\kappa \bar{r})$ indicates the ZLB saddle point. The black line with crosses is the ZLB condition, to the left of which initial conditions are such that the ZLB binds in period t . The solid red line is the stable manifold corresponding to the ZLB saddle point in the case without forward guidance. The region above this line can be interpreted as the immediate recovery region. Under forward guidance, this line becomes flatter and the blue dashed line results. Forward guidance thus has an effect on the immediate recovery region. Assuming that the intensity of choice $\beta = \infty$, households switch from naive expectations to credibility believers in period $t + 2$ for all initial conditions above the blue dotted (forward guidance) and red dash-dotted (no forward guidance) lines, respectively, therefore inducing convergence back to the target steady state. Forward guidance successfully increases this region for which households become credibility believers in period $t + 2$ (indicated by the gray shaded area in panel B), if the credibility of central bank was high enough in the previous periods.

Board) the distribution of expectations clearly becomes bimodal, with a fraction of agents still expecting the same high values of inflation and another fraction expecting lower inflation. In terms of our model, we can interpret this as follows. Before Volcker was appointed the Federal Reserve System (Fed) had very little credibility and most agents expected inflation to remain at the high values that it had been in the recent past

(they used the naive heuristic). In the following quarters the Fed gained more credibility and an increasing fraction of agents started to believe that Volcker would be able to drive down inflation toward its target level (more agents started to follow the fundamental/credibility heuristic). Furthermore, when in 1982 actual inflation started to decline, the mass on high inflation expectations slowly started to move toward lower inflation. We

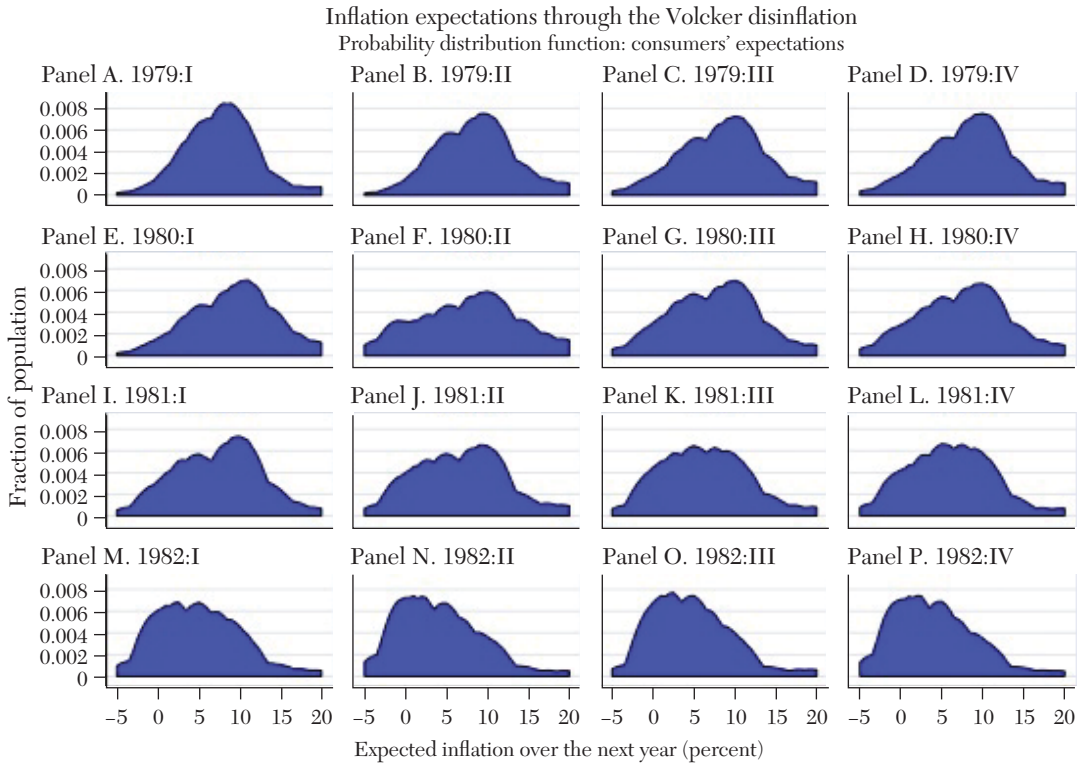


Figure 10. Michigan Survey Data

Note: Bimodal distribution in Michigan Survey data on inflation expectations from 1979 to 1982 (Mankiw et al. 2004).

can interpret this as backward looking, naive agents believing that lower observed inflation would also mean lower inflation in the future. Toward the end of the sample in 1982 both heuristics thus predict lower inflation, consistent with the distribution of the survey data.

Branch (2004, 2007) fit a heuristic switching model with, amongst others, a naive heuristic and a fundamentalistic VAR heuristic to data from the Michigan Survey of Consumer Attitudes and Behavior. Both of these papers find clear evidence of switching between heuristics based on past performance. Branch (2004) furthermore finds that

both our heuristics are present in the survey data, and Branch (2007) shows that the heuristic switching model better fits the survey data than a static sticky information model.²¹

Pfajfar and Žakelj (2016, 2018) and Assenza et al. (2021) show that in their laboratory experiments in the NK framework,

²¹Lux (2009) estimated the parameters of a dynamic opinion formation process with social interactions based on survey data on business expectations (sentiment index data). Madeira and Zafar (2015) use the Michigan Survey data to estimate a learning model of inflation expectations, allowing for heterogeneous use of both private information and lifetime inflation experience.

expectations of subjects can quite accurately (both qualitatively and quantitatively) be described as switching between simple heterogeneous forecasting heuristics based on their relative past performance; see subsection 3.2 for further details.

Cornea-Madeira, Hommes, and Massaro (2019) estimate an NKPC assuming expectations are formed by a heuristic switching model with fundamentalist and naive agents. Fundamentalists here make use of the forward-looking relation between inflation and marginal cost and use a VAR model to make inflation forecasts. Cornea-Madeira, Hommes, and Massaro (2019) find that their model fits the data well and that the endogenous mechanism of switching between the two heuristics based on past performance is supported by the data.

The NKPC with inflation driven by marginal costs and fundamentalists versus naive expectations is given by:

$$(65) \quad \pi_t = \delta \left(n_{f,t} E_t^f \pi_{t+1} + (1 - n_{f,t}) E_t^n \pi_{t+1} \right) + \gamma mc_t + \xi_t,$$

where the fundamental and naive forecasts are given by

$$(66) \quad E_t^f \pi_{t+1} = \gamma e' (I - \delta A)^{-1} A Z_t,$$

$$(67) \quad E_t^n \pi_{t+1} = \pi_{t-1}.$$

Fundamentalists estimate a VAR model and their fraction is given by

$$(68) \quad n_{f,t} = \frac{1}{1 + \exp \left(\beta \left(\frac{FE_{t-1}^f - FE_{t-1}^n}{FE_{t-1}^f + FE_{t-1}^n} \right) \right)},$$

$$(69) \quad FE_{t-1}^i = \sum_{k=1}^K |E_{t-k-1}^i \pi_{t-k} - \pi_{t-k}|,$$

with $i = f, n$.

Figure 11 shows that the one-period-ahead forecast of the heuristics switching model closely matches US inflation. Figure 12 shows the evolution of the fraction of fundamentalists $n_{f,t}$, the distance of actual inflation from the fundamental, the forecast errors of the naive heuristic, and a scatterplot of the fraction of fundamentalists against the relative forecast error of the naive rule. It is clear from this figure that the fraction of fundamentalists varies considerably over time with periods in which it is close to 0.5 and other phases in which it is close to either one of the extremes 0 or 1. For example, immediately after the oil crisis of 1973, the proportion of fundamentalists drops almost to 0 and the naive forecasting rule dominates. Soon after the difference between inflation and fundamental value reaches its peak in 1974:IV, the estimated weight of the forward-looking component shoots back up to about 0.6. During the second oil crisis, inflation was far above the fundamental, causing more agents to adopt a simple backward-looking rule to forecast inflation. Our findings suggest that, in reaction to large shocks pushing inflation away from the fundamental, a large share of agents adopt random walk beliefs causing self-fulfilling high inflation persistence. This result is in line with the analysis of Branch and Evans (2017) showing that innovations to inflation can lead agents adaptively learning in the economy to temporarily believe that inflation follows a random walk. The two-type switching model can also explain a liquidity trap with highly persistent low inflation, as shown in Hommes and Lustenhouwer (2019a) (see subsection 2.3.3).

Cornea-Madeira, Hommes, and Massaro (2019) also estimate the model using survey of professional forecasters (SPF) data. Figure 13 shows the SPF forecasts together with the heuristics switching expectations estimated from the model. Interestingly the behavioral heuristics switching model (HSM) forecasts better match the high peaks of




Figure 11. Actual versus Predicted Inflation

inflation than do the professional forecasters. Cornea-Madeira, Hommes, and Massaro (2019) also estimated structural breaks and figure 13 (right plot) shows that these structural breaks match well with the endogenous switching between fundamentalists and naive expectations for inflation and SPF.

3. *Experimental Macroeconomics*

In the previous section several behavioral models of learning have been discussed. Which learning model is the most relevant? A concern with learning theory is that, “anything goes,” that is, for any equilibrium one can design a suitable theory or learning algorithm that makes that equilibrium

stable under learning. As suggested by Lucas (1986) (see the earlier quote in the introduction), laboratory experiments can provide empirical evidence about the collective behavior and coordination process of a population of adaptive learning agents. Laboratory experiments thus provide a complementary tool to test the empirical relevance of different theories of learning.

Macroeconomists have long been skeptical about the relevance of lab experiments for macro. Nevertheless, aggregate market behavior has been studied in the lab since the early days of experimental economics. Smith (1962), for example, showed the stability of equilibrium in double auction laboratory




Figure 12. Estimation Results of 2-Type Switching Model to US inflation Data

Notes: Panel A: time series of the fraction of fundamentalists $n_{f,t}$. Panel B: distance between actual and fundamental inflation. Panel C: distance between inflation and naive forecast. Bottom: Scatterplot $n_{f,t}$ versus relative forecast error naive rule.

markets. At the other extreme, the seminal contribution of Smith, Suchanek, and Williams (1988) and many follow-up papers have shown the emergence and prevalence of bubbles and crashes in experimental asset markets; see, for example, the survey by Noussair and Tucker (2013).

Following Duffy (2016, p. 4) a macroeconomic experiment can be defined as “one that tests the predictions of a macroeconomic model or its assumptions.” Experimental macroeconomics thus provides complementary tools to falsify or test macro

models in a controlled environment. This seems particularly relevant for behavioral macroeconomics, because lab experiments can provide empirical guidance for which individual decision rules are most relevant in behavioral macro modeling. A rapidly increasing interest in experimental macroeconomics is witnessed, for example, in the extensive handbook survey chapters of Duffy (2016), Arifovic and Duffy (2018), and Mauersberger and Nagel (2018).

The most important characteristic of a macro experiment is that it is a group

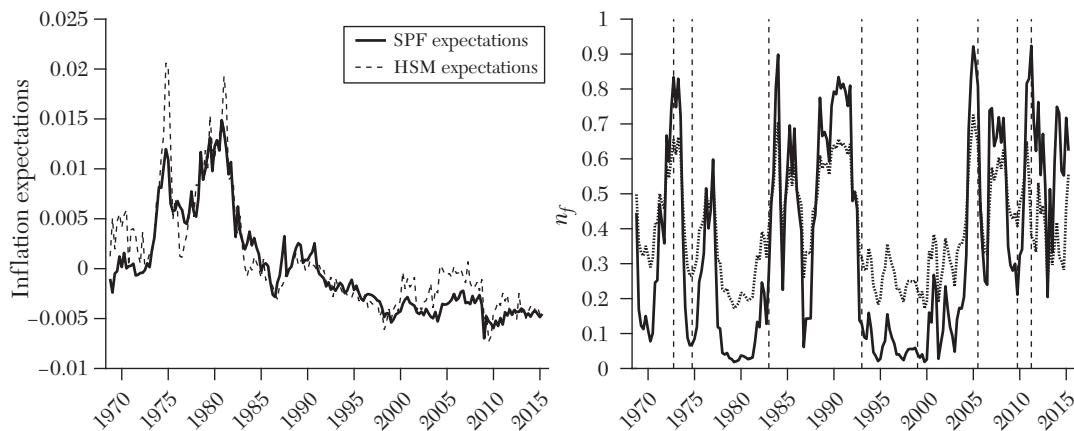


Figure 13. Estimation of 2-Type Switching Model on SPF Data

Notes: SPF forecasts versus HSM expectations and estimated structural breaks with fractions of fundamentalists for inflation and SPF. The right plot shows that the professional forecasters switch slower than the behavioral HSM, which fits inflation data better.

experiment, where individual decisions affect aggregate outcomes, which then feed back into individual behavior, etc. In a controlled macro experiment one can therefore simultaneously test individual (micro) decision rules, their interactions, and the emergent aggregate (macro) behavior they cocreate. An important issue is the size of the group, which, for a macro experiment is often taken to be between five and ten to distinguish it from game theoretic experiments of group size two or three. With group size larger than three individual strategic behavior already becomes very complex. See for example the collection of papers in Duffy (2014) and the discussions about macro experiments therein.

In this section we discuss broadly two types of macro experiments. First, learning-to-forecast (LtF) experiments focusing on how individuals form expectations and how these expectations aggregate. LtF experiments provide a laboratory test of the expectations hypothesis in a given

macro environment. The second type of experiments are policy experiments studying the effectiveness of different policy scenarios in macroeconomic environments. Here we will focus on monetary and fiscal policies.

3.1 Learning-to-Forecast Experiments

Learning-to-forecast (LtF) experiments were pioneered by Marimon and Sunder (1993, 1994, 1995) and Marimon, Spear, and Sunder (1993). Their design is tailor-made to study individual expectations and learning in standard macroeconomic frameworks. In LtF experiments subjects' only task is to forecast future variables, while all other agents' actions (consumption, production, investment, trading, etc.) are computerized typically following rational assumptions from an underlying benchmark macro model or theory. LtF experiments are thus analogous to most of the learning literature in that agents are only boundedly rational in

terms of forecasting, not optimizing.²² LtF experiments may be viewed as an empirical test of coordination of expectations within a given modeling framework. We will also briefly discuss differences with the *learning-to-optimize* experiments, where subjects directly engage in quantity decisions (e.g. consumption, production, trade, etc.).²³ Hommes (2011); Hommes (2013, chap. 8); and Assenza et al. (2014) provide earlier surveys on learning-to-forecast (LtF) experiments.

LtF experiments provide insights into the following questions:

- How do individuals form expectations and learn and adapt their behavior?
- What is the aggregate outcome or emergent macro behavior of individual interactions and learning?
- Will coordination occur or will heterogeneity persist?
- Will adaptive behavior enforce convergence to REE or can nonrational equilibria arise at the macro level?

3.1.1 Asset Pricing Experiments

In the asset pricing LtF experiments in Hommes et al. (2005) there are two assets: a risk free asset paying a fixed rate of return r ; and a risky asset, with price p_t , paying an uncertain dividend y_t . The asset market is populated by six large pension funds and a small fraction of fundamentalist robot traders. Six subjects are forecast advisers to each of the pension funds. Subjects' only task is to forecast the price p_{t+1} of the risky

asset for 50 periods and, based on this forecast, the pension fund then computes how much to invest in the risky asset according to a standard mean-variance demand function. The fundamentalist robot trader always predicts the *fundamental price* p^f and trades based upon this prediction. The realized asset price in the experiment is derived by market clearing and given by:

$$(70) \quad p_t = \frac{1}{1+r} \left((1-n_t) \bar{p}_{t+1}^e + n_t p^f + \bar{y} + \varepsilon_t \right),$$

where $\bar{p}_{t+1}^e = (\sum_{h=1}^6 p_{h,t+1}^e) / 6$ is the average two-period-ahead price forecast, $p^f = \bar{y} / r$ is the fundamental price, and ε_t are small shocks. Subjects do *not* know the underlying law of motion (70), but they do know the mean-dividend \bar{y} and the interest rate r , so they could use these to compute the fundamental price and use it in their forecast. The fraction n_t in (70) is the share of computerized fundamental robot traders, increasing as the price moves away from the fundamental benchmark according to

$$(71) \quad n_t = 1 - \exp\left(-\frac{1}{200} |p_{t-1} - p^f|\right).$$

The fundamental trader thus acts as a “far from equilibrium” stabilizing force in the market, adding negative feedback when the asset price becomes overvalued. The negative feedback becomes stronger the more price moves away from fundamental. The overall expectations feedback system (70) has positive feedback, but the positive feedback becomes less strong (i.e., stronger mean-reverting) when price moves away from fundamental value.

Figure 14 shows time series of prices, individual predictions, and forecasting errors in three different groups with a robot trader. A striking feature of aggregate price behavior

²²An exception is Evans and McGough (2018), where agents must also learn to optimize.

²³Bernasconi and Kirchkamp (2000) combine the learning-to-forecast and learning-to-optimize designs in an overlapping generations (OLG) model experiment; see also Arifovic, Hommes, and Salle (2019) who study learning-to-forecast and learning-to-optimize experiments in a complex OLG environment with infinitely many equilibria.

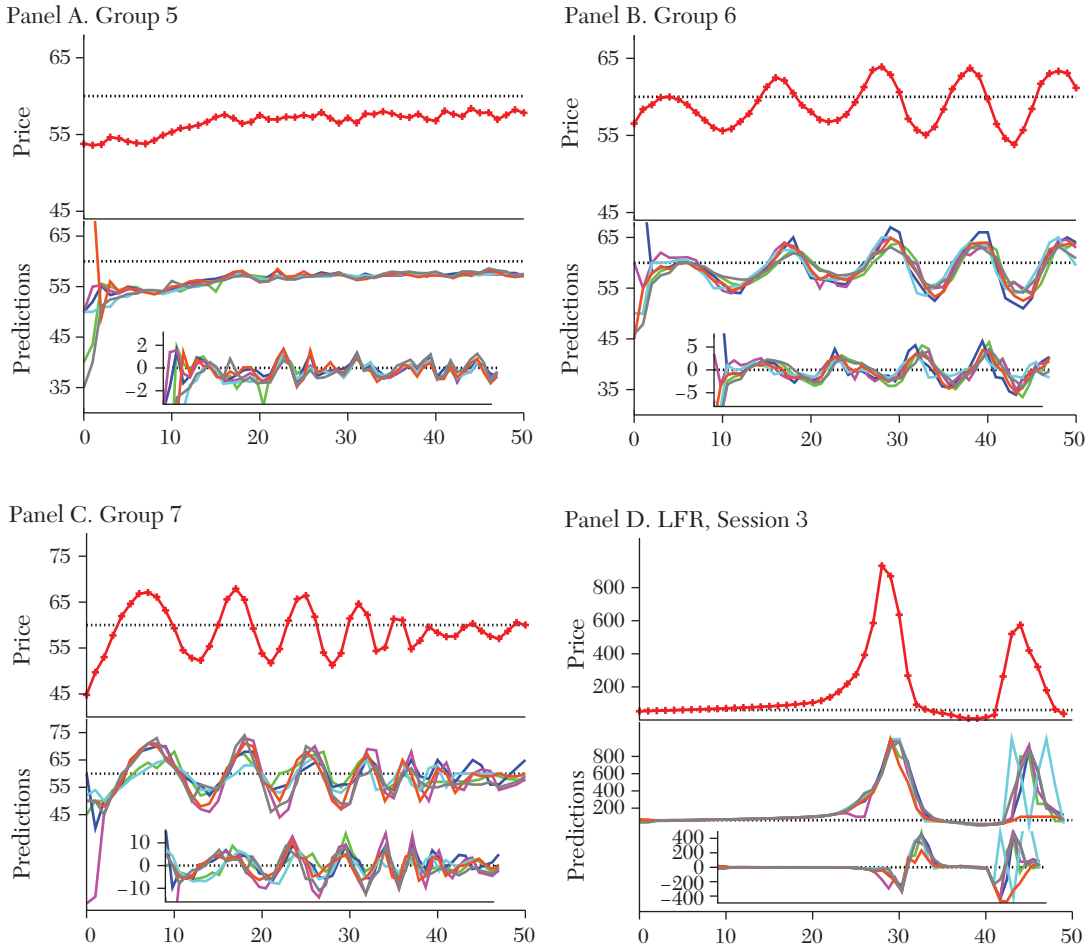


Figure 14. Asset Market Experiments

Notes: Asset pricing experiments (Hommes et al. 2005) realized market prices (upper part each panel), six individual predictions (middle part each panel), and individual errors (bottom part of each panel). Three asset markets with robot traders (panels A, B, and C) and one asset market without robot traders (bottom right). Prices do not converge to the RE fundamental benchmark 60, but rather fluctuate. In the market without fundamental robot trader (panel D, as in Hommes et al. 2008) a long-lasting bubble arises. Individual expectations coordinate on almost self-fulfilling equilibria.

is that three different qualitative patterns emerge. The price in group 5 converges slowly and almost monotonically to the fundamental price level $p^f = 60$. In group 6 persistent oscillations are observed during

the entire experiment, while in group 7 prices fluctuate but the amplitude is decreasing.

A second striking result is that in all groups participants were able to coordinate their forecasts. The forecasts, as shown in

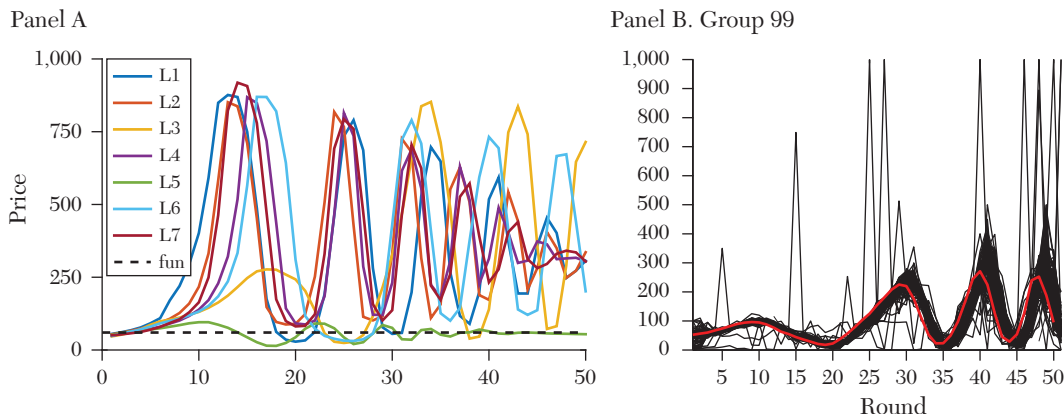


Figure 15. Large Group Experiments

Notes: Panel A: large bubbles in groups of 23–32 (Bao et al. 2020). Panel B: coordination of individual expectations on a large bubble in a group of 100 subjects by coupling two laboratories in Amsterdam and Valencia each with 50 subjects (Hommes, Kopányi-Peuker, and Sonnemans forthcoming).

the lower parts of the panels, are dispersed in the first periods but then, within three–five periods, move close to each other. The coordination of individual forecasts has been achieved in the absence of any communication between subjects, other than through the realized market price, and without any knowledge of past and present predictions of *other* participants.

The fourth group in figure 14 shows a time series of prices, in a market *without* fundamental traders (Hommes et al. 2008). In the absence of a far-from-equilibrium stabilizing force due to negative feedback from the fundamental robot traders, a long-lasting asset price bubble occurs with asset prices rising above 900, that is, more than 15 times the fundamental price, before reaching an exogenously imposed upper-bound of 1,000 and a subsequent market crash. Similar large and long-lasting bubbles have been observed in larger groups of 20–32 (Bao et al. 2020) and even for groups up to 100 (Hommes, Kopányi-Peuker, and Sonnemans forthcoming; see figure 15). Coordination

on bubbles is thus robust against group size.²⁴

These asset market laboratory experiments exhibit a strong degree of coordination on price fluctuations. Markets do not converge to the unique perfectly self-fulfilling RE fundamental price of 60, but rather fluctuate persistently and exhibit expectations driven bubbles and crashes. Subjects therefore do not coordinate on the unique RE equilibrium, but rather coordinate on an almost self-fulfilling equilibrium with temporary bubbles, where forecasting errors are relatively small.

3.1.2 Coordination Failures in Positive Feedback Systems

In his classical paper introducing rational expectations Muth already noted that a

²⁴Recent work in Kopányi-Peuker and Weber (forthcoming) and Hennequin (2019) shows that bubbles are also robust with respect to experience, as bubbles repeatedly appear when subjects gain experience in repeated markets.

crucial feature for aggregation of individual expectation is whether the deviations of individual expectations from the rational forecast are *correlated* or not. To quote Muth (1961, p. 321, emphasis added):

Allowing for cross-sectional differences in expectations is a simple matter, because their *aggregate affect is negligible* as long as the deviation from the rational forecast for an individual firm is *not strongly correlated with those of the others*. Modifications are necessary only if the correlation of the errors is large and depends systematically on other explanatory variables.

Laboratory experiments are well suited to study correlation of individual expectations in a controlled environment. It turns out that the type of expectations feedback, positive or negative, is crucial. In the case of positive (negative) feedback, an increase (decrease) of the average forecast, causes the realized market price to rise (fall). Positive and negative feedback are closely related to strategic complementarity and substitutability (Haltiwanger and Waldman 1985; Fehr and Tyran 2001, 2005, 2008).²⁵ Under negative feedback, when the average forecast goes up, realized price goes down, so it is better to go against the majority (strategic substitutability). Under positive feedback, in contrast, when average forecast goes up, the realized price goes up, so it is better to go with the majority (strategic complementarity). Positive feedback seems particularly relevant in speculative asset

²⁵Fehr and Tyran (2001, 2005, 2008) study the role of money illusion and show that differences in the strategic environments (complementarity versus substitutability) has an impact on individual rationality and aggregate outcomes. Fehr and Tyran (2008) study the adjustment of nominal prices after an anticipated nominal shock in a price setting game with positive (complements) and negative (substitutes) reaction curves, and find much faster convergence in the case of substitutes. They argue that differences in the stickiness of price expectations are key for the understanding of these differences in aggregate outcomes.

markets. If many agents expect the price of an asset to rise they will start buying the asset, aggregate demand will increase and so, by the law of supply and demand, the asset price will increase. High price expectations then become self-fulfilling leading to high realized asset prices. In markets where the role of speculative demand is less important, for example, in markets for non-storable commodities, negative feedback may play a more prominent role. For example in a supply-driven commodity market, if many producers expect future prices to be high they will increase production which, according to the law of supply and demand, will lead to a lower realized market price.

Heemeijer et al. (2009) investigate how the expectations feedback structure affects individual forecasting behavior and aggregate market outcomes by considering market environments that *only* differ in the sign of the expectations feedback, but are equivalent along all other dimensions. The realized price is a linear map of the average of the individual price forecasts $p_{i,t}^e$ of six subjects. The (unknown) price-generating rules in the *negative* and *positive* feedback systems were respectively:²⁶

²⁶This LtF experiment may be viewed as a repeated guessing game as in Nagel (1995), where subjects predict a number between 0 and 100 and the winner is she whose guess is closest to two-thirds of the average. This is a repeated Keynes' beauty contest, where one has to guess the average opinion of other subjects. The Nash equilibrium of the guessing game is 0, but in the laboratory experiment first- and second-order rationality (where the subject guesses $2/3 \cdot 50$ respectively $(2/3)^2 \cdot 50$) are most common. The key difference here is that subjects do not know the best response function, but only have qualitative information about the market. Such limited knowledge seems particularly relevant in macroeconomic systems. Sutan and Willinger (2009) investigate a new version of the beauty contest games (BCG) in which players' actions are strategic substitutes (negative feedback) versus strategic compliments (positive feedback) and find that chosen numbers are closer to rational play in the case of strategic substitutes. See Mauersberger and Nagel (2018) for an

$$(72) \quad p_t = 60 - \frac{20}{21} \left[\left(\sum_{i=1}^6 \frac{1}{6} p_{i,t}^e \right) - 60 \right] + \epsilon_t,$$

(negative feedback)

$$(73) \quad p_t = 60 + \frac{20}{21} \left[\left(\sum_{i=1}^6 \frac{1}{6} p_{i,t}^e \right) - 60 \right] + \epsilon_t,$$

(positive feedback)

where ϵ_t is an exogenous random shock to the pricing rule. The *only* difference between (72) and (73) is the sign of the slope of the linear map, $20/21 \approx +0.95$ resp. $-20/21 \approx -0.95$.²⁷ Heemeijer et al. (2009) consider positive and negative feedback systems with small IID shocks $\epsilon_t \sim N(0, 0.25)$. Negative feedback markets are rather stable and converge quickly to equilibrium, while positive feedback markets are rather unstable and fluctuate around equilibrium, as illustrated in figure 16.

Here we focus on the experiments of Bao et al. (2012), with large permanent shocks to the fundamental price level. More precisely, these shocks have been chosen such that, both in the negative and positive feedback treatments, the fundamental equilibrium price p_t^* changes over time according to:

$$(74) \quad \begin{aligned} p_t^* &= 56, & 0 \leq t \leq 21, \\ p_t^* &= 41, & 22 \leq t \leq 43, \\ p_t^* &= 62, & 44 \leq t \leq 65. \end{aligned}$$

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 17 shows for positive and negative feedback the average price behavior (top panels), realized prices in all groups (middle panels), and an example of individual forecasts in a positive as well as a negative feedback group (bottom panels). 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 18 reveals some other striking features of aggregate price behavior and individual forecasts. The left panel shows the time variation of the median distance to the RE benchmark price over all (eight) groups in both treatments. For the negative feedback treatment, after each large shock the distance spikes, but converges quickly back (within five–six periods) to almost 0. In the positive feedback treatment after each shock the distance to the RE benchmark shows a similar spike, but falls back only slowly and does not converge to 0. The right panel shows how the *degree of heterogeneity*, that is, the median standard deviation of individual forecasts, changes over time. For the positive feedback treatment after each large shock heterogeneity decreases very quickly and converges to (almost) 0 within 3–4 periods. Under positive feedback, individuals thus coordinate expectations quickly, but they all coordinate on the “wrong,” that is, a non-RE price. In the negative feedback treatment heterogeneity is more persistent for about ten periods after each large shock. Persistent heterogeneity stabilizes

extensive overview of experimental coordination games and their importance for macroeconomics.

²⁷In both treatments, the absolute value of the slopes is 0.95, implying in both cases that the feedback system is stable under naive expectations.

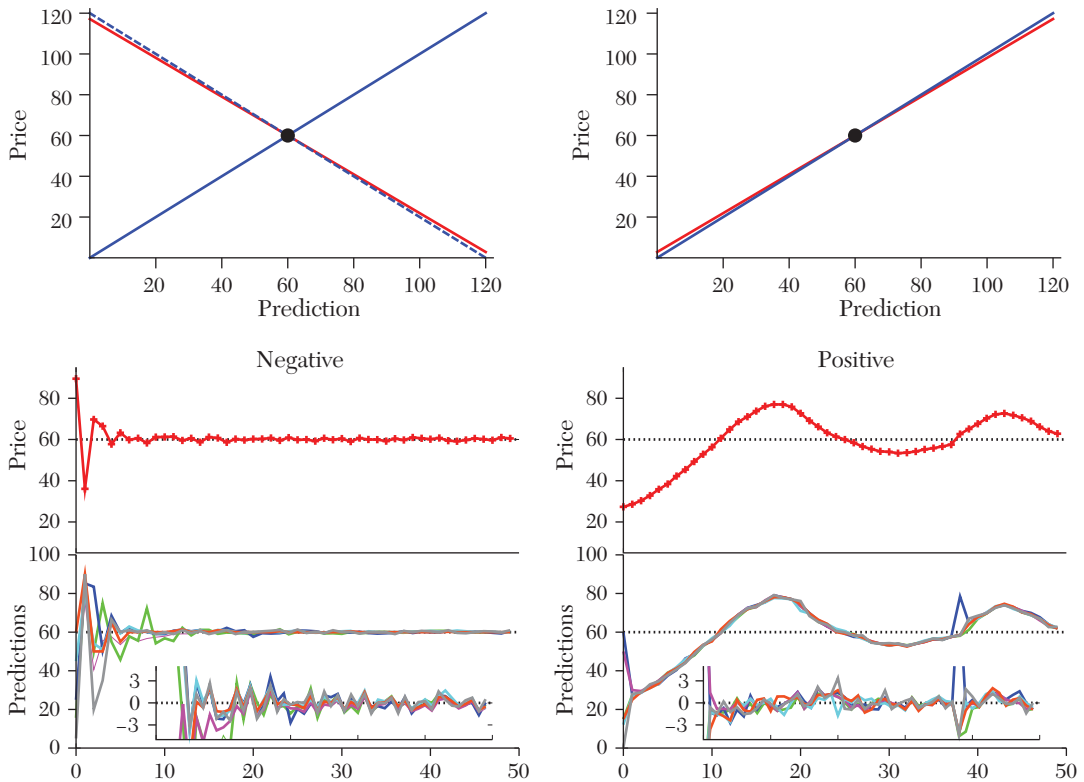


Figure 16. Positive versus Negative Feedback Experiments.

Notes: Laboratory experiments with negative feedback (left panels) and positive feedback (right panels). Upper panels show negative feedback map, with RE steady state 60 clearly visible, and positive feedback map where the unique RE equilibrium is 60, but every other point is *almost* an equilibrium. Other panels show realized market prices (middle panels), six individual predictions (bottom panels), and individual errors (small bottom panels). In the negative expectations feedback market (left panels) the realized price quickly converges to the RE benchmark 60. In positive feedback markets a coordination failure arises and individuals coordinate on the “wrong” price forecast and as a result the realized market price persistently deviates from the RE benchmark 60.

price fluctuations and after convergence of the price to its RE fundamental individual expectations coordinate on the correct RE price.

One may summarize these results in saying that in the positive feedback treatment individuals quickly coordinate on a common prediction, but that coordination on the “wrong” non-fundamental price occurs. As a result price behavior is very different from

the perfect, homogeneous rational expectations equilibrium price. On the other hand, in the negative feedback treatment coordination is much slower, heterogeneity is more persistent, but price convergence is quick.

Stated differently, positive feedback systems are characterized by quick and persistent *coordination failures*, while negative feedback markets are characterized by slow coordination, more persistent heterogeneity,


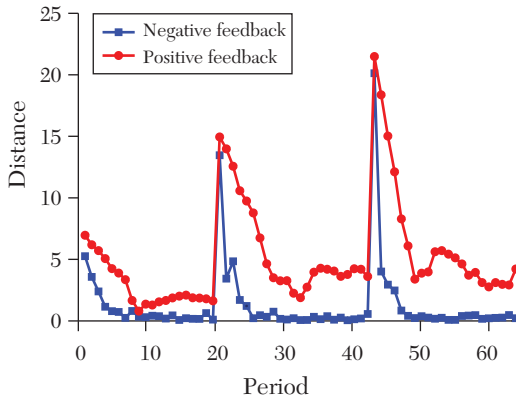


Figure 17. Experiments with Large Shocks

Notes: Positive feedback (left panels) and negative feedback (right panels) experiments with large shocks. Top panels: the average realized price averaged over all eight groups. Middle panels: the market prices for eight different groups. Bottom panels: predictions of six individuals in group P8 (left) and group N8 (right) plotted together with fundamental price (dotted lines). The positive feedback markets are characterized by persistent coordination failures.

Panel A. The median of the distance between the market price and RE



Panel B. The median of the standard deviation

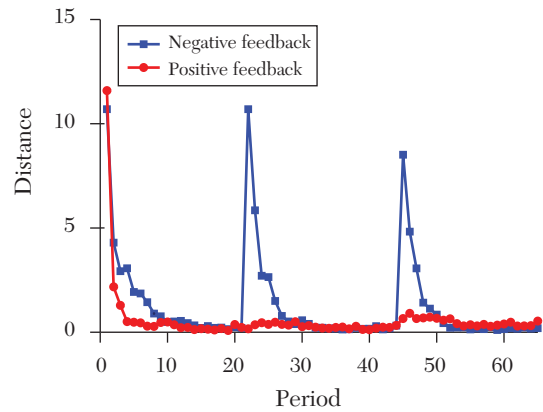


Figure 18. Distance to Fundamental and Individual Forecast Heterogeneity

Notes: Positive/negative feedback markets with large shocks. These plots illustrate price discovery (panel A) and coordination of individual expectations (panel B). The left panel shows the median absolute distance to RE fundamental price, while the right panel shows the median standard deviation of individual predictions. In positive feedback markets coordination is quick, but on the “wrong,” that is, non-RE, price.

and quick price discovery. Notice also that under positive feedback, coordination on a non-RE-fundamental price is *almost self-fulfilling*, with small individual forecasting errors. The positive feedback market is thus characterized by coordination on almost self-fulfilling equilibria with prices very different from the perfectly rational self-fulfilling equilibrium.²⁸ Similar results have been obtained in laboratory experiments in other market settings, including a new Keynesian macro framework (Adam 2007, Pfajfar and Zakelj 2016, Assenza et al. 2014) and in a Lucas asset pricing model (Asparouhova et al. 2016).

²⁸Wagener (2014) uses the same experimental data and shows weak individual rationality (i.e., unbiased forecast errors without autocorrelations) for both the negative and positive feedback treatments, but strong rationality (i.e., prices converge to the homogeneous REE price) only under negative feedback.

3.1.3 Heuristics Switching Model

The fact that qualitatively different aggregate outcomes arise suggests that *heterogeneous expectations* must play a key role to explain these experimental data. Anufriev and Hommes (2012b), extending the model of Brock and Hommes (1997), fitted a behavioral heuristics switching model (HSM) to explain individual forecasting as well as aggregate price behavior.

Agents choose from a number of simple *forecasting heuristics*. The forecasting heuristics are similar to those obtained from estimating linear models on individual forecasting experimental data. *Evolutionary selection* or *performance-based reinforcement learning* based upon relative performance disciplines the individual choice of heuristics. Hence, the impact of each of the rules is evolving over time and agents tend to switch to more

successful rules. The four forecasting heuristics are:

$$(75) \quad p_{1,t+1}^e = 0.65p_{t-1} + 0.35p_{1,t}^e, \quad (\text{ADA})$$

$$(76) \quad p_{2,t+1}^e = p_{t-1} + 0.4(p_{t-1} - p_{t-2}), \quad (\text{WTR})$$

$$(77) \quad p_{3,t+1}^e = p_{t-1} + 1.3(p_{t-1} - p_{t-2}), \quad (\text{STR})$$

$$(78) \quad p_{4,t+1}^e = \frac{p_{t-1}^{av} + p_{t-1}}{2} + (p_{t-1} - p_{t-2}), \quad (\text{LAA})$$

were $p_{t-1}^{av} = \sum_{j=0}^{t-1} p_j$ is the sample average of past prices. *Adaptive expectations* (ADA) predicts that the price is a weighted average of the last observed price p_{t-1} and the last price forecast p_t^e . The *trend-following rules* extrapolate the last price change, either with a weak-trend rule (WTR) or with a strong-trend rule (STR) parameter. The fourth rule is a *learning anchor and adjustment* (LAA) rule (Tversky and Kahneman 1974), extrapolating a price change from a more flexible anchor.

These four rules have been chosen in line with the estimated linear rules for individual forecasts and correspond to the different types of behavior observed in the experiment. ADA leads to monotonic convergence to the fundamental price. The WTR also leads to convergence, possibly after some overshooting and/or oscillatory behavior. The STR leads to instability and a large asset price bubble. Finally, the LAA rule leads to oscillatory behavior. While the strong-trend rule can not predict a price reversal, as it always predicts a price trend to continue, the LAA rule predicts a price reversal when the price moves too far away from the fundamental equilibrium price due to its more flexible anchor that gives 50 percent weight to the average price p_{t-1}^{av} , which may be seen as a proxy for the fundamental equilibrium level.

The fractions of the four forecasting heuristics are time-varying and evolve according

to a discrete choice model with *asynchronous updating*:

$$(79) \quad n_{i,t} = \delta n_{i,t-1} + (1 - \delta) \frac{\exp(\beta U_{i,t-1})}{\sum_{i=1}^4 \exp(\beta U_{i,t-1})}.$$

The fitness or performance measure of forecasting heuristic i is based upon quadratic forecasting errors, consistent with the earnings in the experiments:

$$(80) \quad U_{i,t-1} = -(p_{t-1} - p_{i,t-1}^e)^2 + \eta U_{i,t-2},$$

where $\eta \in [0, 1]$ measures the strength of the agents' *memory*. In the special case $\delta = 0$, (79) reduces to the *the discrete choice model* with synchronous updating; δ represents inertia in switching as subjects change strategies only occasionally. The parameter $\beta \geq 0$ represents the intensity of choice measuring how sensitive individuals are to differences in strategy performance.²⁹

Figure 19 compares the experimental data with the *one-step-ahead predictions* made by the HSM. The one-step-ahead simulations use exactly the same information available to participants in the experiments. The one-period-ahead forecasts easily follow the different patterns in aggregate price behavior in all groups. The second and bottom panels show the corresponding fractions of the four heuristics for each group. In different groups, different heuristics are dominating the market after starting off from an equal distribution.

In the monotonically converging group, the impact of the different rules stays more or less equal, although the impact of adaptive expectations gradually increases and slightly dominates the other rules in the last 25 periods. In the oscillatory group the LAA rule dominates the market from the start and its impact increases to about 90 percent toward the end of the experiment. For the group with

²⁹In the simulations below the parameters are fixed at the benchmark values $\beta = 0.4$, $\eta = 0.7$, $\delta = 0.9$, as in Anufriev and Hommes (2012b).




Figure 19. HSM Model Fitted to Laboratory Experiments

Notes: Simulated prices in laboratory experiments in different groups (red) with corresponding one-step-ahead predictions of the HSM (blue), predictions and forecasting errors (inner frames) of four heuristics and time series of fractions of each of the four heuristics—adaptive expectations (ADA, purple), weak trend followers (WTR, black), strong trend followers (STR, blue), and learning anchor and adjustment heuristic (LAA, red). Panels A–F correspond to three groups with robot traders; bottom panels correspond to group without robot trader and large bubble (panels G, and J) and negative (panels H and K) and positive (panels I and L) feedback groups. In the negative feedback market the ADA rule dominates and enforces quick convergence to the RE fundamental price 60. In the positive expectations feedback market, the strong (STR) and the weak (WTR) trend-following rules perform well and reinforce price oscillations. In all positive feedback groups individual expectations coordinate on a non-RE almost self-fulfilling equilibrium.

the dampened oscillations, one-step-ahead forecast produces rich evolutionary selection dynamics, with three different phases where the STR, the LAA, and the ADA heuristics subsequently dominate. The STR dominates during the initial phase of a strong trend in prices, but starts declining after it misses the first turning point of the trend. The LAA does a better job in predicting the trend reversal, because of its more slowly time-varying anchor and its impact starts increasing. The LAA takes the lead in the second phase of the experiment, with oscillating prices, and its share increases to almost 90 percent after 35 periods. But the oscillations slowly dampen and therefore, after period 35, the impact of adaptive expectations, which has been the worst-performing rule until that point, starts increasing and adaptive expectations dominates the group in the last nine periods. In the asset market without a fundamental trader, subjects coordinate on the strong trend-following strategy, thus explaining the large bubble in the experiment.

The HSM also matches aggregate price behavior in both the negative and positive feedback experiments very well (see the two bottom right panels in figure 19). The time series of the fractions of the different forecasting heuristics provide an intuitive explanation of how individual learning leads to different aggregate price behavior. In the negative feedback treatment, the adaptive expectations strategy performs best and within 20 periods it captures more than 90 percent of the market, thus enforcing convergence toward the RE fundamental equilibrium price. In contrast, in the positive feedback treatment the strong and weak trend-following rules dominate the market, amplifying price fluctuations. The difference in aggregate behavior is thus explained by the fact that *trend-following rules are successful in a positive feedback environment* reinforcing price oscillations and persistent deviations from the fundamental equilibrium

benchmark price, while the trend-following rules are driven out by adaptive expectations in the case of negative feedback (Anufriev and Hommes 2012a). Self-confirming coordination on trend-following rules in a positive expectations feedback environment has an aggregate effect with realized market prices deviating significantly and persistently from the RE benchmark.

3.1.4 Simple Heuristics That Make Us Smart

The HSM provides an intuitive explanation of the laboratory LtF experiments. But an important question remains unanswered: where exactly do the forecasting heuristics with these coefficients come from? Anufriev, Hommes, and Makarewicz (2019) develop a model where agents use a genetic algorithm to optimize the parameters of an anchor and adjustment heuristic. The two-parameter forecasting heuristic is given by

$$(81) \quad p_{i,h,t}^e = \alpha_{i,h,t} p_{t-1} + (1 - \alpha_{i,h,t}) p_{i,t-1}^e + \beta_{i,h,t} (p_{t-1} - p_{t-2}).$$

Here, $\alpha \in [0,1]$ represents the *anchor*, determining how much weight is given to the last observed price versus the last forecast, while β represents the *trend* extrapolation parameter. Each agent i has a set of $H = 20$ heuristics, which are updated over time by a genetic algorithm through reproduction, mutation, and election, with rules being selected with a probability proportional to its relative performance. Anufriev, Hommes, and Makarewicz (2019) show that this GA model outperforms the HSM and adaptive learning benchmarks in explaining and forecasting different laboratory experimental data sets. What makes the GA model work particularly well is the use of an appropriate forecasting heuristic that takes the trend extrapolation in positive feedback systems into account. The LAA heuristic

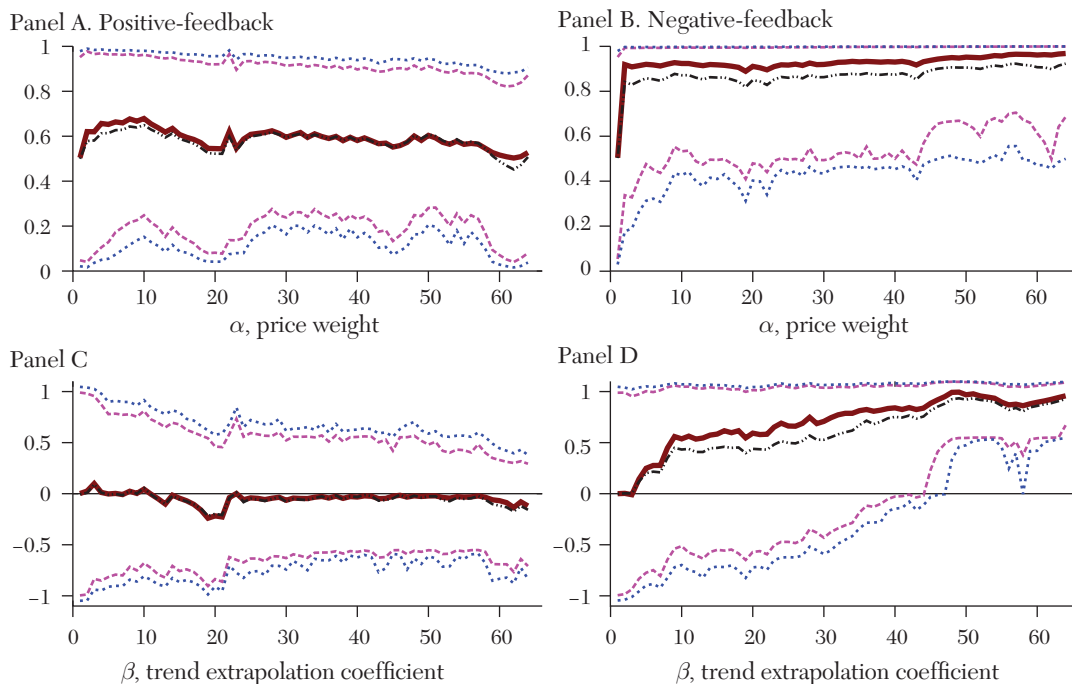


Figure 20. Learning of Anchor and Adjustment Heuristics

Notes: Bao et al. (2012): emerging heuristics in 65-period-ahead MC simulation (1,000 runs). The price weight α (panels A and B) and the trend extrapolation coefficient β (panels C and D) of the chosen heuristic are shown. Red thick line is the median, black dot-dashed line is the mean, blue dotted and purple dashed lines show the 95 percent and 90 percent confidence intervals (CIs), respectively, for the GA learning simulations.

makes the GA learning “smart” in the sense that it fits well with the observed behavior of human subjects in the experiments; cf. Gigerenzer, Todd, and the ABC Research Group (1999).

Figure 20 illustrates which heuristics were learned by the GA agents in the lab experiment of Bao et al. (2012). The figure shows the median and the mean (with 95 percent and 90 percent CI) for 1,000 runs in a Monte Carlo simulation of the price weight α and the trend extrapolation coefficient β , which were selected by the six GA agents. A first observation is that large heterogeneity of individual rules persists, consistent with the

estimated rules in Bao et al. (2012). Secondly, there are clear differences between the two treatments. Under positive feedback the median GA agent quickly converges toward an approximate rule

$$(82) \quad p_{i,t+1}^e \approx 0.95p_t + 0.05p_{i,t}^e + 0.9(p_t - p_{t-1}).$$

This median rule is close to a pure trend-following rule (i.e., with anchor p_t). The 95 percent CI for the trend extrapolation coefficient β becomes significantly positive toward the end of the experiment.

Hence, in the positive feedback environment with large shocks, GA agents learn to become strong trend followers.

These results show how agents learn a *par-simonious* two-parameter forecasting heuristic that makes them smart and fits well with human behavior in laboratory experiments. Hommes et al. (2017) apply this GA-model to the NK environment with similar results.

3.1.5 Learning to Optimize

LtF experiments are analogous to most of the adaptive learning literature in that agents are only boundedly rational in terms of forecasting, but are fully rational in terms of optimizing. In LtF experiments these optimization decisions are computerized. Experiments that directly solicit consumption, production, investment, etc., decisions are called learning-to-optimize (LtO) experiments in the literature (Marimon, Spear, and Sunder 1993). In such LtO experiments there are (at least) two degrees of freedom, namely forecasting and quantity decisions. There are a few papers studying whether the results of LtF experiments are robust with respect to an LtO design. Bao, Duffy, and Hommes (2013) compare LtF and LtO experiments under negative feedback in a cobweb model framework and conclude that the coordination and convergence to the stable REE arising in the LtF design is robust in the LtO experiments, but the convergence speed under LtO is slower than under LtF.

Bao, Hommes, and Makarewicz (2017) compare LtF and LtO in positive feedback asset markets. Under positive feedback, LtO does not lead to convergence to the REE, but leads to large and persistent fluctuations around the fundamental benchmark, with large repeated bubbles of similar magnitude (as measured by price-to-fundamental ratios) as the dotcom stock market or the US housing bubble. These bubbles become even larger under LtO than under LtF, suggesting that LtO is even harder than LtF.

These experimental results suggest that we need more behavioral models that not only replace rational expectations by learning, but also replace optimization by simple and more plausible decision heuristics; see subsection 2.1.7 on the anticipated utility approach and the recent survey Branch and McGough (2018). This remains an important area for future work.

3.2 Policy Experiments in the NK Framework

This subsection discusses macro policy experiments. Cornand and Heinemann (2014) give a recent survey on experiments on monetary policy and central banking. Our main focus here is experiments where the interactions between individual expectations and different policy rules play an important role. A central question then is how different policies may affect the coordination process in standard macroeconomic settings. Learning-to-forecast experiments allow us to reproduce a stylized artificial macroeconomy working along the lines of standard macro models used by academic and policy institutions, with the important difference that no a priori assumptions are made regarding expectations, but instead expectations are directly elicited from incentivized human subjects.

In an early LtF experiment with an expectational Phillips curve, Arifovic and Sargent (2003) use an environment to study the time consistency problem of Kydland and Prescott (1977). One subject acts as policy maker, setting inflation up to a random error term, while a group of three, four, or five subjects forecasts the inflation rate. The experimental results show heterogeneity of expectations across subjects and estimation of an adaptive expectations rule indicates that most subjects formed their forecast by heavily overweighting the recent past. From the monetary policy perspective, Arifovic and Sargent's findings show that in 9 out of 12 experimental economies the policy maker

pushes inflation near the Ramsey value for many periods. Moreover, *backsliding*, that is, inflation drifting back toward the Nash value, occurs in 4 out of 12 economies.

Adam (2007) implements a sticky price environment where inflation and output depend on expected inflation. The results show cyclical patterns of inflation around its steady state. Adam finds that in most of the experimental sessions, the average forecast of subjects is well described by a simple AR(1) model consistent with a restricted perception equilibrium.

A number of laboratory experiments have studied the stabilizing effects of a Taylor-type interest rate rule within a new Keynesian framework, for example, Pfajfar and Žakelj (2016, 2018), Kryvtsov and Petersen (2013), and Assenza et al. (2019). These laboratory experiments provide empirical support to the Taylor principle, that a more aggressive interest rate rule can stabilize inflation and output. All these experiments are similar in spirit, but have important differences in the experimental designs.

Pfajfar and Žakelj (2016) focus on expectation formation in an NK environment. They find that for 30–45 percent of subjects it is not possible to reject rationality. Moreover, 20–25 percent of subjects' forecasting strategies are well described by adaptive learning algorithms. The authors also find evidence for simple heuristics. Roughly 25–35 percent of subjects can be described by trend extrapolation rules and an additional 10–15 percent by adaptive expectations or by a sticky information type of model. Pfajfar and Žakelj (2016) also find evidence for switching between forecasting models.

Kryvtsov and Petersen (2013) provide subjects with full information about an exogenous shock process. This setup allows estimating forecasts as a function of the observed shock history, which is then used to quantify the contribution of expectations to macroeconomic stabilization via counterfactual

analysis. They show that a model with a weak form of adaptive expectations, attributing a significant weight on $t - 1$ realizations of inflation and the output gap, fits best both the magnitude and the timing of aggregate fluctuations observed in the experiment.

In the NKPC framework Assenza et al. (2021) show that simple first-order forecasting heuristics describe individual forecasting behavior well. A behavioral HSM, with agents switching between ADA, STR, WTR, and LAA rules based upon relative performance provides a good description of individual and aggregate behavior. Using the HSM, Assenza et al. (2019) provide a behavioral explanation of why the Taylor principle works. A more aggressive Taylor interest rate rule adds negative feedback to the macro system, thus weakening the positive feedback and making coordination on destabilizing trend-following behavior less likely, thus dampening inflation and output fluctuations. More aggressive monetary policy thus influences the coordination and self-organization of the NK macro system in favor of more stabilizing forecasting heuristics such as adaptive expectations. Hommes, Massaro, and Weber (2019) study the stabilizing effect of a Taylor rule that targets both inflation and the output gap and show that output stabilization can lead to less volatility in inflation. These results are in line with a behavioral HSM.

Survey data on expectations have received much attention in recent years; see the extensive overview of Coibion, Gorodnichenko, and Kamdar (2018). Cornand and Hubert (2020) discuss the external validity of expectations inflation forecasts in the lab. They conclude that overall inflation forecast data from lab experiments and surveys share common features: lagged inflation positively affects the determination of inflation expectations, forecast errors are comparably large and autocorrelated, and forecast errors and forecast revisions are predictable from past information, suggesting the presence of

some form of bounded rationality or information imperfections.

The ZLB and Recovery from Liquidity Traps.—Two recent experimental papers (Arifovic and Petersen 2015; and Hommes, Massaro, and Salle 2019) focus on the effectiveness of monetary and fiscal policies in the new Keynesian framework with a ZLB on the interest rate. In both experiments liquidity traps with inflation and output jointly falling arise in the laboratory, either triggered by shocks to fundamentals as in Arifovic and Petersen (2015), or purely expectations driven as in Hommes, Massaro, and Salle (2019). A fiscal switching rule can recover the economy from the liquidity trap, although recovery may be very slow.

The design in Hommes, Massaro, and Salle (2019) is based on the nonlinear new Keynesian model with multiple RE equilibria, a target steady state, and a ZLB saddle steady state (Evans, Guse, and Honkapohja 2008, see subsection 2.1.5 of this paper). Recall from figure 4 that for initial states below the stable manifold of the ZLB steady state the economy under adaptive learning falls into a liquidity trap in the form of a deflationary spiral. The purpose of the lab experiment is to test whether such deflationary spirals occur in the lab, and whether monetary and/or fiscal policy can recover the economy when expectations are formed by subjects in the lab. They design four (2×2) treatments, with two different policies and two different expectations treatments. The policies are (M) an aggressive monetary policy only, that cuts the interest rate if inflation falls below a threshold of 1.6 percent, and (F) a fiscal switching rule such that, when inflation falls below its threshold even with a ZLB, government spending is increased until inflation reaches the threshold again. The two expectations treatments are: (P) a pessimistic expectations treatment (such that the midpoint of the interval of expectations lies

in the unstable region), and (S) an optimistic expectations treatment (with the midpoint of expectations in the stable region) followed by a number of negative expectational shocks (in periods 8, 9, and 10). Figure 21 shows the results of the four treatments:

- in the MP treatment, five out of seven economies fall into a deflationary spiral;
- in the MS treatment, after the expectations shocks four out of seven economies fall into a deflationary spiral;
- in the FP treatment, the fiscal switching rule keeps all seven economies stable, converging to the target;
- in the FS treatment, despite the expectational shocks all seven economies recover and converge to target.

Figure 22 illustrates the behavior of the MP and MS treatments in the inflation-output phase space. In the MP treatment the position of the initial average expectations exactly determines the outcome: only the five economies whose initial expectations are in the unstable region, below the stable manifold of the ZLB saddle steady state, fall into liquidity traps. For the MS treatment the behavior is similar: only the economies for which, after the bad expectational shocks, the state moves into the unstable region below the stable saddle path of the ZLB fall into deflationary spirals. These laboratory outcomes are strikingly similar to what the theory of adaptive learning predicts. Notice that the lab outcomes are very different from what RE predicts, namely that the system will jump to the target steady state or to the ZLB. In this standard nonlinear NK macro framework the lab experiments are more in line with adaptive learning.

Laboratory experiments should become a complimentary testing bed for new policies. After all, one may raise the question that if a policy does not work in a simple laboratory macro environment, why would

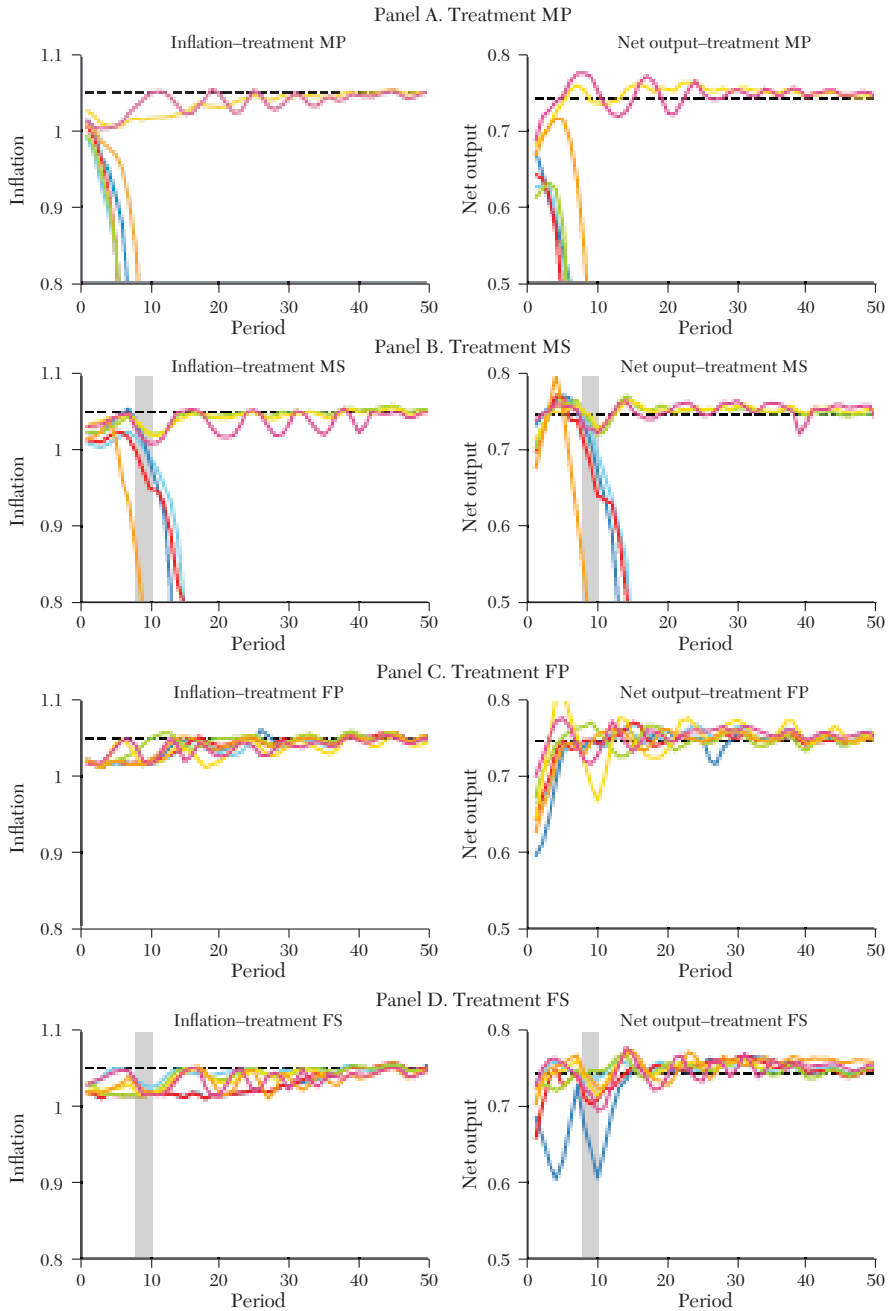


Figure 21. Inflation and Output in ZLB Experiments

Notes: Overview of experimental results of the four treatments, seven groups each, in Hommes, Massaro, and Salle (2019). Left panels: realized inflation. Right panels: realized net output. Dashed lines depict targeted equilibrium levels. Shaded areas indicate expectational bad news shocks.

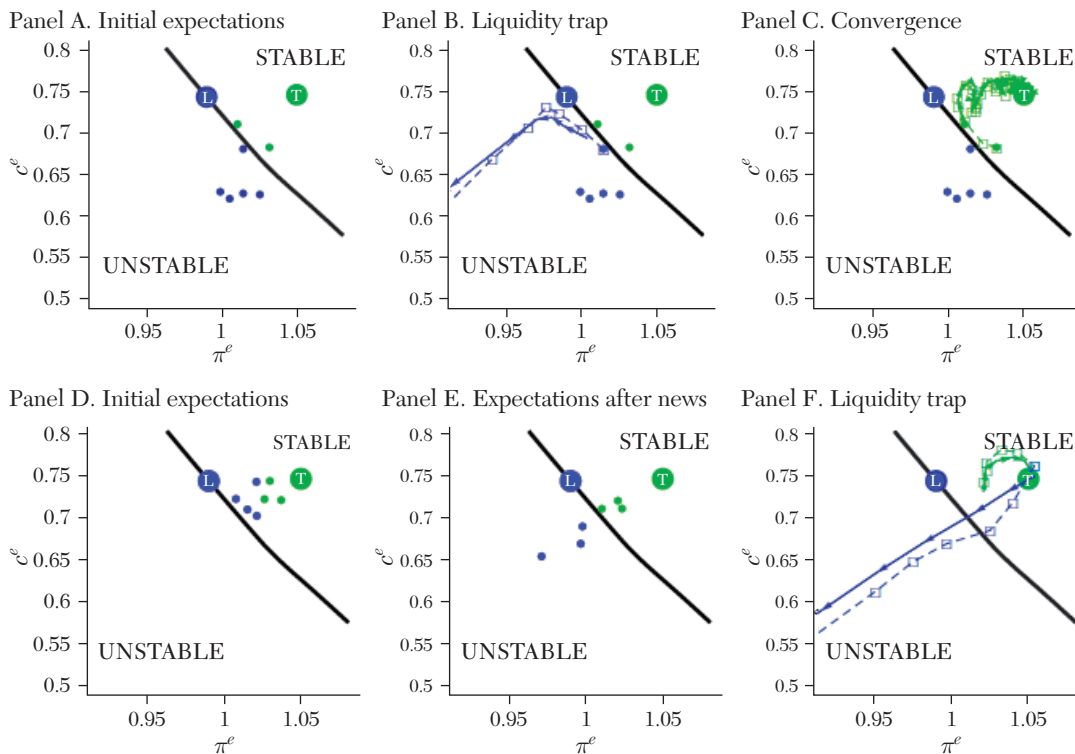


Figure 22. Learning in ZLB Experiments

Notes: Inflation-output dynamics in the seven laboratory groups in the MP (top panels) and the MS (bottom panels) treatments in Hommes, Massaro, and Salle (2019). The black curve is the stable manifold of the ZLB saddle steady state L and marks the boundary between the stable region converging to the target steady state T and the unstable inflationary spirals. The dots refer to the initial average expectations of each group. The predictions of the theory of adaptive learning exactly coincide with the group behavior in the lab.

it work in reality? Testing the effectiveness of policies in the lab thus seems a necessary step to validate policies. The effectiveness of unconventional macroeconomic policies have recently been studied in laboratory experiments. The results on the effect of forward guidance on economic stability in new Keynesian learning-to-forecast experiments are mixed. Cornand and M'baye (2018a,b) find that communication of the central bank's inflation target can reduce the volatility of the economy in normal times,

while Arifovic and Petersen (2015) find that it does not provide a stabilizing anchor in crisis times or in a liquidity trap. Mokhtarzadeh and Petersen (2020) find that providing the economy with the central bank's projections for inflation and the output gap stabilizes the economy, while Kryvtsov and Petersen (2013) find that providing the expected future interest rate path diminishes the effectiveness of monetary policy in stabilizing the economy. Ahrens, Lustenhouwer, and Tettamanzi (2017) show that the central

bank can significantly manage market expectations through forward guidance and that this management strongly supports monetary policy in stabilizing the economy and reducing forecast errors. Moreover, strategically optimistic forward guidance drastically reduces the probability of a deflationary spiral after strong negative shocks to the economy. However, pessimistic forward guidance announcements after negative shocks can instead initiate coordination on a deflationary spiral. Although the credibility of the central banks' forecasts depends on both the central banks' forecasts errors and on credibility in earlier periods, they do not find evidence that a central bank with a better forecasting track record is better able to mitigate recessions than a central bank with less credible past forecasts.

Penalver et al. (forthcoming) report the results of a repeated laboratory experiment in which a central bank buys bonds for cash in a quantitative easing (QE) operation in an otherwise standard asset market setting. The experiment is designed so that bonds have a constant fundamental value which is not affected by QE under rational expectations. By repeating the same experience three times, they investigate whether participants learn that prices should not rise above the fundamental price in the presence of QE. Some groups do learn the fundamental price but most do not, instead learning to believe that QE boosts bond prices. In future work, more laboratory policy experiments are needed to investigate the robustness of these results and these complimentary laboratory methodologies should become part of the standard tools for policy analysis.

4. Discussion and Future Outlook

Modern macroeconomics is built on elegant and rigorous axiomatic micro-foundations and rational expectations. But are these assumptions empirically relevant?

Laboratory macro-experiments show that coordination of a collection of adaptive agents is better described by a behavioral learning process than by rational expectations equilibrium. In particular, macro systems with strong *positive feedback* (i.e., near-unit root systems) typically exhibit coordination failures in the lab. A collection of individuals does not coordinate on the perfect rational equilibrium, even when it is unique, but rather coordinates on almost self-fulfilling equilibria characterized by correlated trend-following behavior and booms and bust price fluctuations around the rational fundamental benchmark. This behavior has been observed in many lab experiments within different environments, including asset markets and new Keynesian macro frameworks, and such aggregate behavior is robust for larger group sizes up to 100 subjects. What are the policy implications of these empirically and experimentally observed coordination failures?

Before discussing the policy implications of these empirical and experimental results, it is useful to recall some features of what has become the standard and widely used model for policy analysis in central banks, the RE new Keynesian DSGE model. Estimation of DSGE models based on a Bayesian likelihood approach has been pioneered by Smets and Wouters (2003, 2007). Their model has seven endogenous macroeconomic variables, incorporates many types of real and nominal frictions, and uses seven structural exogenous shock processes, all autocorrelated AR(1) or autoregressive moving average (ARMA) (1,1) processes, to match the number of observables. Under these assumptions the model has a *unique, determinate REE saddle-path equilibrium* solution. Smets and Wouters (2003) estimate 19 structural parameters of the model using the Dynare advanced computational software, together with 17 parameters for the structural shocks. The estimated coefficients

of the shock AR(1) processes exhibit highly persistent autocorrelation including three estimated coefficients in the range 0.95–0.97. Many central banks and financial institutions use similar estimated DSGE models as an important input for policy analysis. Yet, coordination on a saddle-path equilibrium seems highly unlikely and has rarely been observed in laboratory macro experiments.³⁰

In order to further discuss rational versus behavioral approaches for policy analysis, consider the following *thought experiment* using the time series of a single observed macro variable obtained from the aggregate price series in the positive feedback laboratory experiments of Heemeijer et al. (2009) (see subsection 3.1.2, figure 16, right panel). The experimenter knows that the price series has been generated by

$$(83) \quad p_t = p^* + \frac{1}{1+r}(\bar{p}_t^e - p^*) + \epsilon_t,$$

where p^* ($= \bar{y}/r = 60$) is the REE fundamental price, r ($= 0.05$) is the risk free interest rate, \bar{p}_t^e is the average forecast of a group of six individuals, and ϵ_t is an IID exogenous stochastic process. A standard REE full information approach would assume that

³⁰A simple example of a saddle point steady state arises in the ZLB lab experiment in Hommes, Massaro, and Salle (2019), as discussed in subsection 3.2. None of the 28 groups of 6 subjects were able to coordinate expectations on the ZLB steady state, nor did any group converge to a neighborhood of the ZLB steady state. In a hyperinflation model with two RE steady states, Marimon and Sunder (1993) find coordination on the low-inflation determinate REE steady state, which is consistent with adaptive learning. In an NK experiment, Adam (2007) finds coordination on the determinate REE steady state in some sessions, but coordination on a restricted perception equilibrium in other sessions. Pfajfar and Žakelj (2016, 2018) and Assenza et al. (2019) find in NK experiments that the Taylor principle is not sufficient to enforce coordination on the unique determinate REE steady state, but almost self-fulfilling fluctuations in output and inflation may arise when monetary policy satisfies the Taylor principle, but is not aggressive enough.

expectations are rational, i.e., coincide with the fundamental steady-state equilibrium p^* , and that prices are driven by exogenous autocorrelated shocks ϵ_t , say by a stochastic AR(1) process. Such an REE model driven by a persistent AR(1) process would give a very reasonable fit to the observed aggregate time series. However, from the individual laboratory forecasting data we know that expectations are very different from homogeneous rational expectations and fundamental price, but rather subjects coordinated on a simple trend-following rule. The behavioral HSM, as discussed in subsection 3.1.3, buffeted with small IID shocks explains both individual forecasting behavior and observed macro fluctuations in prices. Both models, the REE model driven by (large) autocorrelated shocks and the behavioral heterogeneous expectations model driven by small IID shocks, give a good description of aggregate data. Both models, however, have very different policy implications about how the macro system might respond to policy parameters, such as the interest rate r . In the REE model price fluctuations are driven by autocorrelated exogenous shocks, and the policy implication would be that the interest rate only affects the equilibrium price level $p^* = \bar{y}/r$, but does *not* affect the price volatility. In contrast, in the behavioral model the interest rate affects both the price level and the price volatility, because an increase of the interest rate leads to stronger *mean reversion* of the price process. An increase of the interest rate weakens the positive feedback of the system (83) and thus makes coordination on trend-following behavior less likely and the macro system more stable. This simple example is an illustration of the fact that behavioral models may have very different policy recommendations than RE models. The effect upon the steady state price levels is the same, but in behavioral models policy parameters typically have a strong effect upon the *strength of the*

mean-reversion of the system. In behavioral models policy often affects the (in)stability and convergence speed of the model.

How general are these results? Hommes (2013) stresses an important generic feature of (higher-dimensional) near-unit root macro systems. Many macro-finance models, such as asset pricing and new Keynesian models, are near-unit root systems. Since any linear system with a unit root has a continuum (say a line) of steady states, any near-unit root system (i.e., whose Jacobian has an eigenvalue close to +1) has a continuum of *almost self-fulfilling equilibria*. What the laboratory macro experiments show is that in such an environment, coordination on the unique perfectly self-fulfilling rational equilibrium is unlikely and at best only very slow, but coordination of expectations on almost self-fulfilling equilibria fluctuating around the RE fundamental benchmark is likely to emerge. Coordination failures similar to those observed in the simple laboratory experiments are therefore likely to be generic for near-unit root macro systems. Policy analysis should take these almost self-fulfilling equilibria into account. This observation may be seen as a widening of the solution to the Lucas critique. Policy analysis should not only be based on the perfect rational solution, but take these empirically relevant almost self-fulfilling learning equilibria into account.³¹ The laboratory macro experiments also suggest an immediate way to stabilize the economy: policy can stabilize the economy by adding negative feedback to the system, or equivalently, by weakening the positive feedback.

In recent learning-to-forecast experiments of the housing market, Bao and Hommes

(2019) studied the role of negative feedback policies in the housing market. There is positive feedback from speculation and negative feedback from housing construction. The equilibrium housing price is determined by:

$$(84) \quad p_t = \frac{1}{1+r}(\bar{p}_{h,t+1}^e + \bar{y} - c\bar{p}_{i,t+1}^e) + \nu_t,$$

where $\bar{p}_{h,t+1}^e$ is the average forecast made by the speculators, $\bar{p}_{i,t+1}^e$ the average forecast by the constructors, r ($= 0.05$) is the interest rate, \bar{y} is the mean housing rent, and $\nu_t \sim N(0,1)$ represents small demand or supply shocks. As can be seen from (84), the housing price will increase when the average price prediction $\bar{p}_{h,t+1}^e$ made by the speculators goes up and, because of the negative feedback, decrease when the average price prediction $\bar{p}_{i,t+1}^e$ by the constructors goes up. Therefore the housing market exhibits *positive expectations feedback* from the speculative investors, and *negative expectation feedback* from the constructors. The overall feedback strength is determined by the eigenvalue $\lambda = (1-c)/(1+r)$. Bao and Hommes (2019) consider three different treatments: strong positive feedback with low housing supply elasticity ($c = 0$; eigenvalue $\lambda = 0.95$), medium positive feedback with intermediate supply elasticity ($c = 0.1$; eigenvalue $\lambda = 0.85$) and weak positive feedback with high supply elasticity ($c = 0.25$; eigenvalue $\lambda = 0.71$). Figure 23 shows the simulated housing prices by the HSM model against the experimental housing prices in the three treatments. The simulated prices fit the experimental data well. The fractions of the different forecasting heuristics show different patterns in the three different treatments. A typical market in treatment without housing construction and only speculation is dominated by the strong-trend rule, which leads to large bubbles and unstable price fluctuations. A typical market in the medium treatment with some

³¹This relates to recent work on robust policy design as it explicitly seeks to design policies that are robust to deviation from rational expectations. Adam and Woodford (2018), for instance, find that such robustness considerations make it optimal for monetary policy to condition on asset prices (housing prices in their setup).

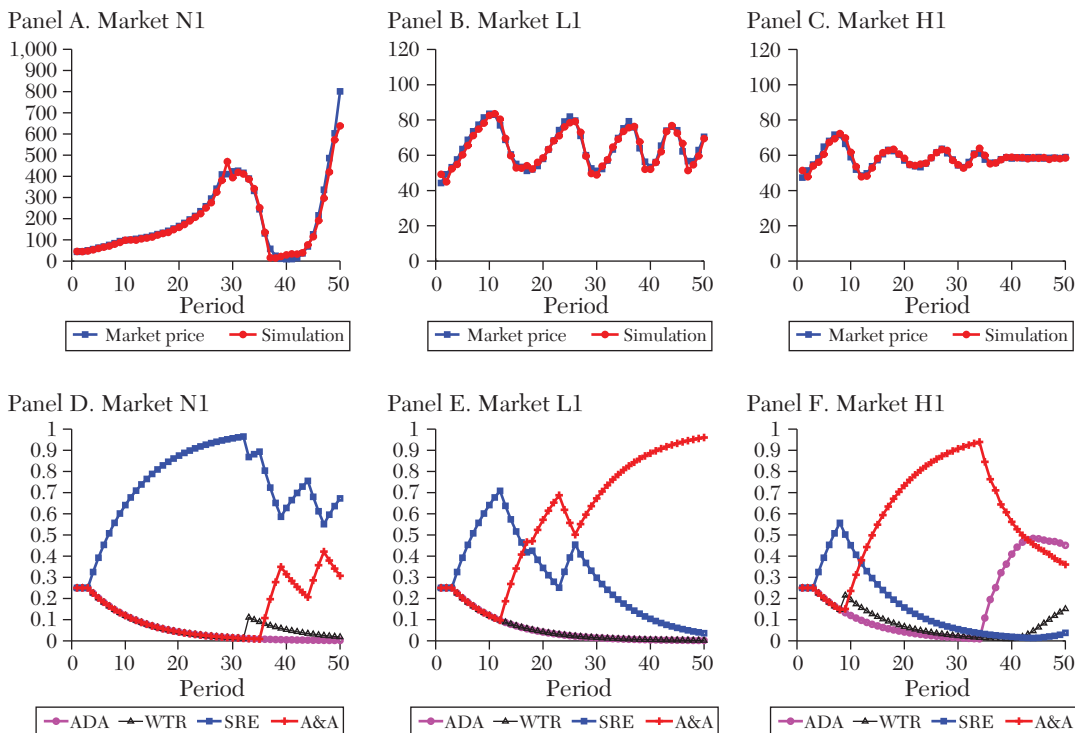


Figure 23. Negative Feedback Policies May Stabilize Housing Prices

Notes: Housing prices (panels A–C) and fractions of the strategies in the HSM. When housing price elasticity is weak there is strong positive feedback and coordination of expectations on trend-following strategies leads to large housing price bubbles. In the intermediate case, bubbles do not arise, but coordination on the LAA rule leads to price oscillations around the fundamental equilibrium. With stronger negative feedback from housing supply trend-following strategies do not survive competition and coordination on ADA stabilizes market prices.

housing construction is initially dominated by the strong-trend rule, but after the reversal of the price trend the LAA rule increases its share and becomes dominating in later periods, which leads to persistent price oscillations. Finally, in the treatment with larger housing supply the market is firstly dominated by the LAA rule, but after period 30 the adaptive rule becomes more popular toward the end of the experiment, which eventually leads to dampening of the oscillations and convergence to the fundamental

price. The HSM thus provides simple and intuitive explanations of this experiment. The large housing bubbles are explained by coordination on a strong trend-following rule; the oscillations in the medium treatment are explained by coordination on an LAA rule, and the stable price behavior in the last treatment is explained by coordination on adaptive expectations. Negative feedback policies in the form of more housing construction thus prevent the strong trend-following strategy to survive and favor

stabilizing adaptive behavior. Policy changes thus affect the self-organization and learning process in the market in favor of stabilizing adaptive expectations in line with the Lucas critique.

Which behavioral models and which form of learning to use then for policy analysis? Parsimonious learning rules and simple forecasting heuristics are natural candidates to start with. The real economy is too complex to fully understand and agents use simple decision heuristics. This is nicely illustrated by the example of the dog and the Frisbee (Haldane 2012). It does not require a PhD in physics and perfect knowledge of Newton's law of motion to catch a Frisbee. A dog can do the job by following a simple adaptive heuristic—run at a speed so that the angle of gaze to the Frisbee remains roughly constant—and humans are likely to follow a similar heuristic when catching a Frisbee on the beach.³² Many simple models of learning have been developed in recent years where agents use simple rules, for example, an optimal AR(1) rule,³³ or choose from a number

³²Haldane's dog and the Frisbee example bears some analogy with Friedman's famous "as if" hypothesis of a billiard player making his shots as if he knew the complicated mathematical formulas that would give the optimum directions of travel, could estimate accurately by eye the angles, etc., describing the location of the balls, could make lightning calculations from the formulas, and could then make the balls travel in the direction indicated by the formulas (Friedman 1953). There are some important differences however. Friedman's example applies to a one-shot game, where with perfect knowledge about the mathematical equations the billiard player plays as if he takes an optimal shot. Haldane's dog and the Frisbee example corresponds to a complex evolving system—the Frisbee in the air—and how a simple adaptive heuristic—keeping the angle of gaze roughly constant—describes the adaptive behavior. Importantly, this adaptive heuristic does not assume perfect knowledge of the underlying laws of motion, but simply adapts to observations of the system.

³³Examples in which DSGE models have taken first steps into the research direction proposed here include Slobodyan and Wouters (2012b, a), who introduce adaptive learning of an AR(2) rule into the Smets–Wouters DSGE model, and Hommes et al. (2021), who consider the NK DSGE model where agents learn an optimal AR(1) forecasting rule for all endogenous variables. The adaptive

of simple forecasting heuristics based upon their relative performance. These behavioral models fit well with empirical data and lab experiments. Future work should focus on parsimonious learning models and include almost self-fulfilling equilibria in the analysis. An empirical micro-foundation, using laboratory experiments, survey data, and other micro decision data, should play a key role in developing behavioral agent-based macro models, stylized as well as computationally more advanced, for more realistic policy analysis in the near future.

REFERENCES

- Adam, Klaus. 2007. "Experimental Evidence on the Persistence of Output and Inflation." *Economic Journal* 117 (520): 603–36.
- Adam, Klaus, Johannes Beutel, Albert Marcet, and Sebastian Merkel. 2015. "Can a Financial Transaction Tax Prevent Stock Price Booms?" *Journal of Monetary Economics* 76 (S): S90–109.
- Adam, Klaus, and Albert Marcet. 2011. "Internal Rationality, Imperfect Market Knowledge and Asset Prices." *Journal of Economic Theory* 146 (3): 1224–52.
- Adam, Klaus, Albert Marcet, and Johannes Beutel. 2017. "Stock Price Booms and Expected Capital Gains." *American Economic Review* 107 (8): 2352–2408.
- Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini. 2016. "Stock Market Volatility and Learning." *Journal of Finance* 71 (1): 33–82.
- Adam, Klaus, and Michael Woodford. 2018. "Leaning against Housing Prices as Robustly Optimal Monetary Policy." NBER Working Paper 24629.
- Adam, Klaus and Michael Woodford. 2012. "Robustly Optimal Monetary Policy in a Microfounded New Keynesian Model." *Journal of Monetary Economics* 59 (5), 468–87.
- Ahrens, Steffen, Joep Lustenhouwer, and Michele Tettamanzi. 2017. "The Stabilizing Role of Forward Guidance: A Macro Experiment." Unpublished.
- Akerlof, George A. 2002. "Behavioral Macroeconomics and Macroeconomic Behavior." *American Economic Review* 92 (3): 411–33.
- Ambrose, Brent W., Piet Eichholz, and Thies Lindenthal. 2013. "House Prices and Fundamentals: 355 Years of Evidence." *Journal of Money, Credit and Banking* 45 (2–3): 477–91.
- Anderson, P. W. 1972. "More Is Different: Broken Symmetry and the Nature of the Hierarchical Structure

learning model provides a better fit to the data with weaker autocorrelations in the exogenous shocks.

- of Science." *Science* 177 (4047): 393–96.
- Anderson, Philip W., Kenneth J. Arrow, and David Pines, eds. 1988. *The Economy as an Evolving Complex System*. Vol. 5 of *Santa Fe Institute Studies in the Sciences of Complexity*. Boca Raton, FL: Taylor & Francis Group, CRC Press.
- 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.
- Anufriev, Mikhail, Tiziana Assenza, Cars Hommes, and Domenico Massaro. 2013. "Interest Rate Rules and Macroeconomic Stability under Heterogeneous Expectations." *Macroeconomic Dynamics* 17 (8): 1574–1604.
- Anufriev, Mikhail, and Cars Hommes. 2012a. "Evolution of Market Heuristics." Special issue, *Knowledge Engineering Review* 27 (2): 255–71.
- Anufriev, Mikhail, and Cars Hommes. 2012b. "Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments." *American Economic Journal: Microeconomics* 4 (4): 35–64.
- 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.
- Arifovic, Jasmina. 1994. "Genetic Algorithm Learning and the Cobweb Model." *Journal of Economic Dynamics and Control* 18 (1): 3–28.
- Arifovic, Jasmina, James Bullard, and Olena Kostyshyna. 2013. "Social Learning and Monetary Policy Rules." *Economic Journal* 123 (567): 38–76.
- Arifovic, Jasmina, Herbert Dawid, Christophe Deisenberg, and Olena Kostyshyna. 2010. "Learning Benevolent Leadership in a Heterogeneous Economy." *Journal of Economic Dynamics and Control* 34 (9): 1768–90.
- Arifovic, Jasmina, and John Duffy. 2018. "Heterogeneous Agent Modeling: Experimental Evidence." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*. Vol. 4, edited by Cars Hommes and Blake LeBaron, 491–540. Amsterdam: North-Holland.
- Arifovic, Jasmina, Cars Hommes, and Isabelle Salle. 2019. "Learning to Believe in Simple Equilibria in a Complex OLG Economy - Evidence from the Lab." *Journal of Economic Theory* 183: 106–82.
- Arifovic, Jasmina, and John Ledyard. 2011. "A Behavioral Model for Mechanism Design: Individual Evolutionary Learning." *Journal of Economic Behavior & Organization* 78 (3): 374–95.
- Arifovic, Jasmina, and Luba Petersen. 2015. "Escaping Expectations Driven Liquidity Traps: Experimental Evidence." Simon Fraser University, Department of Economics Discussion Paper dp15-03.
- Arifovic, Jasmina, Thomas J. Sargent. 2003. "Laboratory Experiments with an Expectational Phillips Curve." In *Evolution and Procedures in Central Banking*, edited by David E. Altig and Bruce D. Smith, 23–56. Cambridge, UK: Cambridge University Press.
- Arifovic, Jasmina, Stephanie Schmitt-Grohé, and Martín Uribe. 2018. "Learning to Live in a Liquidity Trap." *Journal of Economic Dynamics and Control* 89: 120–36.
- Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler. 1997. "Asset Pricing under Endogenous Expectations in an Artificial Stock Market." In *The Economy as an Evolving Complex System II*, edited by W. Brian Arthur, Steven N. Durlauf, and David A. Lane, 15–44. New York: Taylor & Francis Group, CRC Press.
- Ascari, G., N. Pecora, and A. Spelta. 2018. "Boom and Busts in a Housing Market with Heterogeneous Agents." *Macroeconomic Dynamics* 22 (7): 1808–24.
- Asparouhova, Elena, Peter Bossaerts, Nilanjan Roy, and William Zame. "Lucas' in the Laboratory." *Journal of Finance* 71 (6): 2727–80.
- Assenza, Tiziana, Te Bao, Cars Hommes, Domenico Massaro. 2014. "Experiments on Expectations in Macroeconomics and Finance." In *Experiments in Macroeconomics*. Vol. 17 of *Research in Experimental Economics*, edited by John Duffy, 11–70. Bingley, UK: Emerald Publishing.
- 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*. <https://doi.org/10.1016/j.jmoneco.2019.12.005>.
- Bao, Te, John Duffy, and Cars Hommes. 2013. "Learning, Forecasting and Optimizing: An Experimental Study." *European Economic Review* 61: 186–204.
- Bao, Te, Myrna Hennequin, Cars Hommes, and Domenico Massaro. 2020. "Coordination on Bubbles in Large-Group Asset Pricing Experiments." *Journal of Economic Dynamics and Control* 110: Article 103702.
- Bao, Te, and Cars H. Hommes. 2019. "When Speculators Meet Constructors: Positive versus Negative Feedback in Experimental Housing Markets." *Journal of Economic Dynamics and Control* 107: Article 103730.
- Bao, Te, Cars Hommes, and Tomasz Makarewicz. 2017. "Bubble Formation and (in)Efficient Markets in Learning-to-Forecast and Optimize Experiments." *Economic Journal* 127 (605): F581–609.
- 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.
- Baptista, Rafa, J. Dooyne Farmer, Marc Hinterschweiger, Katie Low, Daniel Tang, and Arzu Uluc. 2016. "Macroprudential Policy in an Agent-Based Model of the UK Housing Market." Bank of England Staff Working Paper 619.
- Battiston, Stefano, et al. 2016. "Complexity Theory and Financial Regulation. Economic Policy Needs Interdisciplinary Network Analysis and Behavioral Modeling." *Science* 351 (6275): 818–19.
- Benhabib, Jess, George W. Evans, and Seppo

- Honkapohja. 2014. "Liquidity Traps and Expectation Dynamics: Fiscal Stimulus or Fiscal Austerity?" *Journal of Economic Dynamics and Control* 45: 220–38.
- Benhabib, Jess, Stephanie Schmitt-Grohé, and Martín Uribe. 2002. "Avoiding Liquidity Traps." *Journal of Political Economy* 110 (3): 535–63.
- Bernasconi, Michèle, and Oliver Kirchkamp. 2000. "Why Do Monetary Policies Matter? An Experimental Study of Saving and Inflation in an Overlapping Generations Model." *Journal of Monetary Economics* 46 (2): 315–43.
- Beshears, John, James J. Choi, Andreas Fuster, David Laibson, and Brigitte C. Madrian. 2013. What Goes Up Must Come Down? Experimental Evidence on Intuitive Forecasting. *American Economic Review* 103 (3): 570–74.
- Blanchard, Olivier. 2014. "Where Danger Lurks." *Finance and Development* 51 (3): 28–32.
- Blanchard, Olivier J., and Mark W. Watson. 1982. "Bubbles, Rational Expectations and Financial Markets." In *Crises in Economic and Financial Structure: Bubbles, Bursts, and Shocks*, edited by Paul Wachtel, 295–316. Lexington, MA: DC Heathand Company.
- Blume, Lawrence E., and David Easley. 1993. "Economic Natural Selection." *Economics Letters* 42 (2–3): 281–89.
- Bolt, Wilko, Maria Demertzis, Cees Diks, Cars Hommes, and Marco van der Leij. 2019. "Identifying Booms and Busts in House Prices under Heterogeneous Expectations." *Journal of Economic Dynamics and Control* 103: 234–59.
- Bookstaber, Richard, and Alan Kirman. 2018. "Modeling a Heterogeneous World." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*, Vol. 4, edited by Cars Hommes and Blake LeBaron, 769–95. Amsterdam: North-Holland.
- Boswijk, H. Peter, Cars H. Hommes, Sebastiano Manzan. 2007. "Behavioral Heterogeneity in Stock Prices." *Journal of Economic Dynamics and Control* 31 (6): 1938–70.
- Branch, William A. 2004. "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations." *Economic Journal* 114 (497): 592–621.
- 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.
- Branch, William A. 2007. "Sticky Information and Model Uncertainty in Survey Data on Inflation Expectations." *Journal of Economic Dynamics and Control* 31 (1): 245–76.
- Branch, William A., and George W. Evans. 2010. "Asset Return Dynamics and Learning." *Review of Financial Studies* 23 (4): 1651–80.
- Branch, William A., and George W. Evans. 2011a. "Learning about Risk and Return: A Simple Model of Bubbles and Crashes." *American Economic Journal: Macroeconomics* 3 (3): 159–91.
- Branch, William A., and George W. Evans. 2011b. "Monetary Policy and Heterogeneous Expectations." *Economic Theory* 47 (2–3): 365–93.
- Branch, William A., and George W. Evans. 2017. "Unstable Inflation Targets." *Journal of Money Credit and Banking* 49 (4): 767–806.
- Branch, William, George Evans, and Bruce McGough. 2013. "Finite Horizon Learning." In *Macroeconomics at the Service of Public Policy*, edited by Thomas J. Sargent and Jouku Vilmunen, 141–63. Oxford: Oxford University Press.
- Branch, William A., and Bruce McGough. 2005. "Consistent Expectations and Misspecification in Stochastic Non-linear Economies." *Journal of Economic Dynamics and Control* 29 (4): 659–76.
- Branch, William A., and Bruce McGough. 2009. "A New Keynesian Model with Heterogeneous Expectations." *Journal of Economic Dynamics and Control* 33 (5): 1036–51.
- Branch, William A., and Bruce McGough. 2010. "Dynamic Predictor Selection in a New Keynesian Model with Heterogeneous Expectations." *Journal of Economic Dynamics and Control* 34 (8): 1492–1508.
- Branch, William A., and Bruce McGough. 2018. "Heterogeneous Expectations and Micro-foundations in Macroeconomics." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*, Vol. 4, edited by Cars Hommes and Blake LeBaron, 3–62. Amsterdam: North-Holland.
- Bray, Margaret. 1982. "Learning, Estimation, and the Stability of Rational Expectations." *Journal of Economic Theory* 26 (2): 318–39.
- Bray, M. M., and N. E. Savin. 1986. "Rational Expectations Equilibria, Learning, and Model Specification." *Econometrica* 54 (5): 1129–60.
- Brock, William A. 1993. "Pathways to Randomness in the Economy: Emergent Nonlinearity and Chaos in Economics and Finance." *Estudios Económicos* 8 (1): 3–55.
- Brock, William A., and Steven N. Durlauf. 2001a. "Interactions-Based Models." In *Handbook of Econometrics*, Vol. 5, edited by James J. Heckman and Edward Leamer, 3297–3380. Amsterdam: North-Holland.
- Brock, William A., and Steven N. Durlauf. 2001b. "Discrete Choice with Social Interactions." *Review of Economic Studies* 68 (2): 235–60.
- Brock, William A., and Cars H. Hommes. 1997. "A Rational Route to Randomness." *Econometrica* 65 (5): 1059–95.
- Brock, William A., and Cars H. Hommes. 1998. "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model." *Journal of Economic Dynamics and Control* 22 (8–9): 1235–74.
- Brock, William A., Cars H. Hommes, and Florian O. O. Wagener. 2005. "Evolutionary Dynamics in Markets with Many Trader Types." *Journal of Mathematical Economics* 41 (1–2): 7–42.
- Brock, W. A., C. H. Hommes, and F. O. O. Wagener. 2009. "More Hedging Instruments May Destabilize Markets." *Journal of Economic Dynamics and*

- Control* 33 (11): 1912–28.
- Bullard, James. 1994. “Learning Equilibria.” *Journal of Economic Theory* 64 (2): 468–85.
- Bullard, James, and John Duffy. 2001. “Learning and Excess Volatility.” *Macroeconomic Dynamics* 5 (2): 272–302.
- Bullard, James, George W. Evans, and Seppo Honkapohja. 2008. “Monetary Policy, Judgment and Near-Rational Exuberance.” *American Economic Review* 98 (3): 1163–77.
- Bullard, James, and Kaushik Mitra. 2002. “Learning about Monetary Policy Rules.” *Journal of Monetary Economics* 49 (6): 1105–29.
- Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2016. “Understanding Booms and Busts in Housing Markets.” *Journal of Political Economy* 124 (4): 1088–1147.
- Camerer, Colin, and Teck Hua Ho. 1999. “Experienced-Weighted Attraction Learning in Normal Form Games.” *Econometrica* 67 (4): 827–74.
- Chiarella, Carl, Xue-Zhong He, and Remco C. J. Zwinkels. 2014. “Heterogeneous Expectations in Asset Pricing: Empirical Evidence from the S&P500.” *Journal of Economic Behavior & Organization* 105: 1–16.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. 2018. “On DSGE Models.” *Journal of Economic Perspectives* 32 (3): 113–40.
- Christiano, Lawrence, Roberto Motto, and Massimo Rostagno. 2010. “Financial Factors in Economic Fluctuations.” European Central Bank Working Paper 1192.
- Clarida, Richard, Jordi Galí, and Mark Gertler. 1999. “The Science of Monetary Policy: A New Keynesian Perspective.” *Journal of Economic Literature* 37 (4): 1661–1707.
- Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar. 2018. “The Formation of Expectations, Inflation, and the Phillips Curve.” *Journal of Economic Literature* 56 (4): 1447–91.
- Cornand, Camille, and Frank Heinemann. 2014. “Experiments on Monetary Policy and Central Banking.” In *Experiments in Macroeconomics*. Vol. 17 of *Research in Experimental Economics*, edited by John Duffy, 167–227. Bingley, UK: Emerald Publishing.
- Cornand, Camille, and Paul Hubert. 2020. “On the External Validity of Experimental Inflation Forecasts: A Comparison with Five Categories of Field Expectations.” *Journal of Economic Dynamics and Control* 110: Article 103746.
- Cornand, Camille, and Cheick Kader M’baye. 2018a. “Band or Point Inflation Targeting? An Experimental Approach.” *Journal of Economic Interaction and Coordination* 13 (2): 283–309.
- Cornand, Camille, and Cheick Kader M’baye. 2018b. “Does Inflation Targeting Matter? An Experimental Investigation.” *Macroeconomic Dynamics* 22 (2): 362–401.
- Cornea-Madeira, Adriana, Cars Hommes, and Domenico Massaro. 2019. “Behavioral Heterogeneity in U.S. Inflation Dynamics.” *Journal of Business & Economic Statistics* 37 (2): 288–300.
- Cúrdia, Vasco, and Michael Woodford. 2010. “Credit Spreads and Monetary Policy.” *Journal of Money, Credit and Banking* 42 (S1): 3–35.
- Cúrdia, Vasco, Michael Woodford. 2016. “Credit Frictions and Optimal Monetary Policy.” *Journal of Monetary Economics* 84: 30–65.
- Dawid, Herbert, and Domenico Delli Gatti. 2018. “Agent-Based Macroeconomics.” In *Handbook of Computational Economics: Heterogeneous Agent Modeling*. Vol. 4, edited by Cars Hommes and Blake LeBaron, 63–156. Amsterdam: North-Holland.
- Day, Richard H., and Weihong Huang. 1990. “Bulls, Bears and Market Sheep.” *Journal of Economic Behavior & Organization* 14 (3): 299–329.
- De Grauwe, Paul. 2011. “Animal Spirits and Monetary Policy.” *Economic Theory* 47 (2–3): 423–57.
- De Grauwe, Paul. 2012. *Lectures on Behavioral Macroeconomics*. Princeton: Princeton University Press.
- Del Negro, Marco, Marc Giannoni, and Christina Patterson. 2012. “The Forward Guidance Puzzle.” Federal Reserve Bank of New York Staff Report 574, revised Dec. 2015.
- DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990a. “Positive Feedback Investment Strategies and Destabilizing Rational Speculation.” *Journal of Finance* 45 (2): 379–95.
- DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990b. “Noise Trader Risk in Financial Markets.” *Journal of Political Economy* 98 (4): 703–38.
- Dieci, Roberto, and Xue-Zhong He. 2018. “Heterogeneous Agent Models in Finance.” In *Handbook of Computational Economics: Heterogeneous Agent Modeling*. Vol. 4, edited by Cars Hommes and Blake LeBaron, 257–328. Amsterdam: North-Holland.
- Dieci, Roberto and Frank Westerhoff. 2013. “Modeling House Price Dynamics with Heterogeneous Speculators.” In *Global Analysis of Dynamic Models in Economics and Finance*, edited by Gian Italo Bischi, Carl Chiarelli, and Iryna Sushko, 35–61. Berlin and Heidelberg: Springer.
- Dieci, Roberto and Frank Westerhoff. 2012. “A Simple Model of a Speculative Housing Market.” *Journal of Evolutionary Economics* 22 (62): 303–29.
- Driscoll, John C., and Steinar Holden. 2014. “Behavioral Economics and Macroeconomic Models.” *Journal of Macroeconomics* 41: 133–47.
- Duffy, John, ed. 2014. *Experiments in Macroeconomics*. Vol. 17 of *Research in Experimental Economics*. Bingley, UK: Emerald Publishing.
- Duffy, John. 2016. “Macroeconomics: A Survey of Laboratory Research.” In *The Handbook of Experimental Economics*. Vol. 2, edited by John H. Kagel and Alvin E. Roth, 1–90. Princeton: Princeton University Press.
- Eichholz, Piet, Ronald Huisman, and Remco C. J. Zwinkels. 2015. “Fundamentals or Trends? A Long-term Perspective on House Prices.” *Applied Economics* 47 (10): 1050–59.

- Erev, Ido, and Alvin E. Roth. 1998. "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88 (4): 848–81.
- Eusepi, Stefano, and Bruce Preston. 2017. "The Science of Monetary Policy: An Imperfect Knowledge Perspective." *Journal of Economic Literature* 56 (1): 3–59.
- Evans, George W., and Bruce McGough. 2018. "Learning to Optimize." Unpublished.
- Evans, George W., Eran Guse, and Seppo Honkapohja. 2008. "Liquidity Traps, Learning and Stagnation." *European Economic Review* 52 (8): 1438–63.
- Evans, George W., and Seppo Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press.
- Fehr, Ernst, and Jean-Robert Tyran. 2001. "Does Money Illusion Matter?" *American Economic Review* 91 (5): 1239–62.
- Fehr, Ernst, and Jean-Robert Tyran. 2005. "Individual Irrationality and Aggregate Outcomes." *Journal of Economic Perspectives* 19 (4): 43–66.
- Fehr, Ernst, and Jean-Robert Tyran. 2008. "Limited Rationality and Strategic Interaction: The Impact of the Strategic Environment on Nominal Inertia." *Econometrica* 76 (2): 353–94.
- Föllmer, Hans. 1974. "Random Economies with Many Interacting Agents." *Journal of Mathematical Economics* 1 (1): 51–62.
- Föllmer, Hans, Ulrich Horst, Alan Kirman. 2005. "Equilibria in Financial Markets with Heterogeneous Agents: A Probabilistic Perspective." *Journal of Mathematical Economics* 41 (1–2): 123–55.
- Franke, Reiner, Frank Westerhoff. 2017. "Taking Stock: A Rigorous Modelling of Animal Spirits in Macroeconomics." *Journal of Economic Surveys* 31 (5): 1152–82.
- Frankel, Jeffrey, and Kenneth A. Froot. 1986. "Understanding the US Dollar in the Eighties: The Expectations of Chartists and Fundamentalists." NBER Working Paper R0957.
- Frankel, Jeffrey A., and Kenneth A. Froot. 1987a. "Short-Term and Long-Term Expectations of the Yen/Dollar Exchange Rate: Evidence from Survey Data." *Journal of the Japanese and International Economies* 1 (3): 249–74.
- Frankel, Jeffrey A., and Kenneth A. Froot. 1987b. "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations." *American Economic Review* 77 (1): 133–53.
- Friedman, Milton. 1953. *Essays in Positive Economics*. Chicago and London: University of Chicago Press.
- Froot, Kenneth A., and Maurice Obstfeld. 1991. "Intrinsic Bubbles: The Case of Stock Prices." *American Economic Review* 81 (5): 1189–1214.
- Fuster, Andreas, David Laibson, and Brock Mendel. 2010. "Natural Expectations and Macroeconomic Fluctuations." *Journal of Economic Perspectives* 24 (4): 67–84.
- Fuster, Andreas, Benjamin Hebert, and David Laibson. 2012. "Natural Expectations, Macroeconomic Dynamics, and Asset Pricing." In *NBER Macroeconomics Annual 2011*, Volume 26, edited by Daron Acemoglu and Michael Woodford, 1–48. Chicago: University of Chicago Press.
- Gabaix, Xavier. 2017. "Behavioral Macroeconomics via Sparse Dynamic Programming." Unpublished.
- Gabaix, Xavier. 2020. "A Behavioral New Keynesian Model." *American Economic Review* 110 (8): 2271–2327.
- Geanakoplos, John et al. 2012. "Getting at Systemic Risk via an Agent-Based Model of the Housing Market." *American Economic Review* 102 (3): 53–58.
- Gertler, Mark, and Peter Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* 58 (1): 17–34.
- Gertler, Mark, and Peter Karadi. 2013. "QE 1 vs. 2 vs. 3...: A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool." *International Journal of Central Banking* 2013–2014 (January): 5–53.
- Gigerenzer, Gerd, Peter M. Todd, and the ABC Research Group. 1999. *Simple Heuristics that Make Us Smart*. New York: Oxford University Press.
- Gilchrist, Simon, Alberto Ortiz, and Egon Zakrajšek. 2009. "Credit Risk and the Macroeconomy: Evidence from an Estimated DSGE Model." Unpublished.
- Goy, Gavin, Cars Hommes, and Kostas Mavromatis. 2020. "Forward Guidance and the Role of Central Bank Credibility under Heterogeneous Beliefs." *Journal of Economic Behavior & Organization*. <https://doi.org/10.1016/j.jebo.2020.06.025>.
- Goyal, Sanjeev. 2018. "Heterogeneity and Networks." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*, Vol. 4, edited by Cars Hommes and Blake LeBaron, 687–712. Amsterdam: North-Holland.
- Grandmont, Jean-Michel. 1985. "On Endogenous Competitive Business Cycles." *Econometrica* 53 (5): 995–1045.
- Grandmont, Jean-Michel. 1998. "Expectations Formation and Stability of Large Socioeconomic Systems." *Econometrica* 66 (4): 741–81.
- Haldane, Andrew G. 2012. "The Dog and the Frisbee." Speech, Federal Reserve Bank of Kansas City's 36th Economic Policy Symposium, Jackson Hole, WY, August 31.
- Haldane, A. G., and A. E. Turrell. 2018. "An Interdisciplinary Model for Macroeconomics." *Oxford Review Economic Policy* 34 (1–2): 219–51.
- Haltiwanger, John, and Michael Waldman. 1985. "Rational Expectations and the Limits of Rationality: An Analysis of Heterogeneity." *American Economic Review* 75 (3): 326–40.
- 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.
- Hennequin, Myrna. 2019. "Experiences and Expectations in Asset Markets: An Experimental Study."

- Unpublished.
- Hommes, Cars H. 2006. "Heterogeneous Agent Models in Economics and Finance." In *Handbook of Computational Economics: Agent-Based Computational Economics*. Vol. 2, edited by Leigh Tesfatsion and Kenneth L. Judd, 1109–86. Amsterdam: North-Holland.
- Hommes, Cars. 2011. "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab." *Journal of Economic Dynamics and Control* 35 (1): 1–24.
- Hommes, Cars. 2013. *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge, UK: Cambridge University Press.
- Hommes, Cars, and Daan in 't Veld. 2017. "Booms, Busts and Behavioural Heterogeneity in Stock Prices." *Journal of Economic Dynamics and Control* 80: 101–24.
- Hommes, Cars, Anita Kopányi-Peuker, and Joep Sonnemans. Forthcoming. "Bubbles, Crashes and Information Contagion in Large-Group Asset Markets." *Experimental Economics*. <https://doi.org/10.1007/s10683-020-09664-w>.
- Hommes, Cars, and Blake LeBaron, eds. 2018. *Handbook of Computational Economics: Heterogeneous Agent Modeling*. Vol. 4. Amsterdam: North-Holland.
- Hommes, Cars, and Joep Lustenhouwer. 2019a. "Inflation Targeting and Liquidity Traps under Endogenous Credibility." *Journal of Monetary Economics* 107: 48–62.
- Hommes, Cars, and Joep Lustenhouwer. 2019b. "Managing Unanchored, Heterogeneous Expectations and Liquidity Traps." *Journal of Economic Dynamics and Control* 101: 1–16.
- Hommes, Cars, Tomasz Makarewicz, Domenico Massaro, and Tom Smits. 2017. "Genetic Algorithm Learning in a New Keynesian Macroeconomic Setup." *Journal of Evolutionary Economics* 27 (5): 1133–55.
- Hommes, Cars, Domenico Massaro, and Isabelle Salle. 2019. "Monetary and Fiscal Policy Design at the Zero-Lower-Bound: Evidence from the Lab." *Economic Inquiry* 57 (2): 1120–40.
- Hommes, Cars, Domenico Massaro, and Matthias Weber. 2019. "Monetary Policy under Behavioral Expectations: Theory and Experiment." *European Economic Review* 118: 193–212.
- Hommes, Cars, Kostas Mavromatis, Tolga Özden, and Mei Zhu. 2021. "Behavioral Learning Equilibria in the New Keynesian Model." DeNederlandscheBank Working Paper 654.
- 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.
- Hommes, Cars, Joep Sonnemans, Jan Tuinstra, and Henk van de Velden. 2008. "Expectations and Bubbles in Asset Pricing Experiments." *Journal of Economic Behavior & Organization* 67 (1): 116–33.
- Hommes, Cars, and Gerhard Sorger. 1998. "Consistent Expectations Equilibria." *Macroeconomic Dynamics* 2 (3): 287–321.
- Hommes, Cars, and Mei Zhu. 2014. "Behavioral Learning Equilibria." *Journal of Economic Theory* 150: 778–814.
- Honkapohja, Seppo, Kaushik Mitra, and George W. Evans. 2013. "Notes on Agents' Behavioral Rules under Adaptive Learning and Studies of Monetary Policy." In *Macroeconomics at the Service of Public Policy*, edited by Thomas J. Sargent and Jouko Vilmunen, 63–79. Oxford: Oxford University Press.
- Iori, Giulia, and Rosario N. Mantegna. 2018. "Empirical Analyses of Networks in Finance." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*. Vol. 4, edited by Cars Hommes and Blake LeBaron, 637–85. Amsterdam: North-Holland.
- Kirman, Alan. 1991. "Epidemics of Opinion and Speculative Bubbles in Financial Markets." In *Money and Financial Markets*, edited by Mark P. Taylor, 354–68. Cambridge, UK: Blackwell.
- Kirman, Alan P. 1992. "What or Whom Does the Representative Individual Represent?" *Journal of Economic Perspectives* 6 (2): 117–36.
- Kirman, Alan. 1993. "Ants, Rationality and Recruitment." *Quarterly Journal of Economics* 108 (1): 137–56.
- Kopányi-Peuker, Anita, and Matthias Weber. Forthcoming. "Experience Does Not Eliminate Bubbles: Experimental Evidence." *Review of Financial Studies*. <https://doi.org/10.1093/rfs/hhaa121>.
- Kouwenberg, Roy, and Remco Zwinkels. 2014. "Forecasting the US Housing Market." *International Journal of Forecasting* 30 (3): 415–25.
- Kryvtsov, Oleksiy, and Luba Petersen. 2013. "Expectations and Monetary Policy: Experimental Evidence." Bank of Canada Staff Working Paper 13-44.
- Kurz, Mordecai, ed. 1997. *Endogenous Economic Fluctuations: Studies in the Theory of Rational Beliefs*. Berlin and New York: Springer.
- Kurz, Mordecai, Giulia Piccillo, and Howei Wu. 2013. "Modeling Diverse Expectations in an Aggregated New Keynesian Model." *Journal of Economic Dynamics and Control* 37 (8): 1403–33.
- Kydland, Finn E., and Ed Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (3): 473–92.
- Lansing, Kevin J. 2009. "Time-Varying U.S. Inflation Dynamics and the New Keynesian Phillips Curve." *Review of Economic Dynamics* 12 (2): 304–26.
- Lansing, Kevin J. 2010. "Rational and Near-Rational Bubbles without Drift." *Homica Journal* 120 (549): 1149–74.
- LeBaron, Blake. 2006. "Agent-Based Computational Finance." In *Handbook of Computational Economics: Agent-Based Computational Economics*. Vol. 2, edited by Leigh Tesfatsion and Kenneth L. Judd, 1187–1233. Amsterdam: North-Holland.
- LeBaron, Blake, W. Brian Arthur, and Richard Palmer. 1999. "Time Series Properties of an Artificial Stock Market." *Journal of Economic Dynamics and Control* 23 (9–10): 1487–1516.
- Lof, Matthijs. 2015. "Rational Speculators, Contrarians

- and Excess Volatility." *Management Science* 61 (8): 1889–1901.
- Lucas, Robert E., Jr. 1986. "Adaptive Behavior and Economic Theory." *Journal of Business* 59 (4, part 2): S401–26.
- Lux, Thomas. 1995. "Herd Behaviour, Bubbles and Crashes." *Economic Journal* 105 (431): 881–96.
- Lux, Thomas. 2009. "Stochastic Behavioral Asset Pricing Models and the Stylized Facts." In *Handbook of Financial Markets: Dynamics and Evolution*, edited by Thorsten Hens and Klaus Reiner Schenk-Hoppé, 161–215. Amsterdam: North-Holland.
- Lux, Thomas, and Michele Marchesi. 1999. "Scaling and Criticality in a Stochastic Multi-agent Model of a Financial Market." *Nature* 397 (6719): 498–500.
- Lux, Thomas, and Remco C. J. Zwinkels. 2018. "Empirical Validation of Agent-Based models." In *Handbook of Computational Economics: Heterogenous Agent Modeling*, Vol. 4, edited by Cars Hommes and Blake LeBaron, 437–88. Amsterdam: Elsevier.
- Madeira, Carlos, and Basit Zafar. 2015. "Heterogeneous Inflation Expectations and Learning." *Journal of Money, Credit and Banking* 47 (5): 867–96.
- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2003. "Disagreement about Inflation Expectations." In *NBER Macroeconomics Annual 2003*, Vol. 18, edited by Mark Gertler and Kenneth Rogoff, 209–47. Cambridge, MA: MIT Press.
- Manski, Charles F., and Daniel McFadden. 1981. "Alternative Estimators and Sample Designs for Discrete Choice Analysis." In *Structural Analysis of Discrete Data*, edited by Charles Manski and Daniel McFadden, 2–50. Cambridge, MA: MIT Press.
- Marcet, Albert, and Thomas J. Sargent. 1989a. "Convergence of Least Squares Learning in Environments with Hidden State Variables and Private Information." *Journal of Political Economy* 97 (6): 1306–22.
- Marcet, Albert, and Thomas J. Sargent. 1989b. "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models." *Journal of Economic Theory* 48 (2): 337–68.
- Marimon, Ramon, Stephen E. Spear, and Shyam Sunder. 1993. "Expectationally Driven Market Volatility: An Experimental Study." *Journal of Economic Theory* 61 (1): 74–103.
- Marimon, Ramon, and Shyam Sunder. 1993. "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence." *Econometrica* 61 (5): 1073–1107.
- Marimon, Ramon, and Shyam Sunder. 1994. "Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach." *Economic Theory* 4 (1): 131–62.
- Marimon, Ramon, and Shyam Sunder. 1995. "Does a Constant Money Growth Rule Help Stabilize Inflation?" *Carnegie-Rochester Conference Series on Public Policy* 43: 111–56.
- Massaro, Domenico. 2013. "Heterogeneous Expectations in Monetary DSGE Models." *Journal of Economic Dynamics and Control* 37 (3): 680–92.
- Mauersberger, Felix. 2017. "Monetary Policy Rules in a Non-rational World: A Macroeconomic Experiment." Unpublished.
- Mauersberger, Felix, and Rosemarie Nagel. 2018. "Levels of Reasoning in Keynesian Beauty Contests: A Generative Framework." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*, Vol. 4, edited by Cars Hommes and Blake LeBaron, 541–634. Amsterdam: North-Holland.
- Milani, Fabio. 2007. "Expectations, Learning and Macroeconomic Persistence." *Journal of Monetary Economics* 54 (7): 2065–82.
- Milani, Fabio. 2011. "Expectations, Shocks and Learning as Drivers of the Business Cycle." *Economic Journal* 121 (552): 379–401.
- Mokhtarzadeh, Fatemeh, and Luba Petersen. 2020. "Coordinating Expectations through Central Bank Projections." *Experimental Economics*. <https://doi.org/10.1007/s10683-020-09684-6>.
- Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica* 29 (3): 315–35.
- Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review* 85 (5): 1313–26.
- Noussair, Charles N., and Steven Tucker. 2013. "Experimental Research on Asset Pricing." *Journal of Economic Surveys* 27 (3): 554–69.
- Penalver, Adrian, Nobuyuki Hanaki, Eizo Akiyama, Yukihiko Fumaki, and Ryuichi Ishikawa. 2018. "An Experimental Analysis of the Effect of Quantitative Easing." Banque de France Working Paper 684.
- Pfajfar, Damjan, and Blaž Žakelj. 2016. "Experimental Evidence on Inflation Expectation Formation." *Journal of Economic Dynamics and Control* 44: 147–68.
- Pfajfar, Damjan, and Blaž Žakelj. 2018. "Inflation Expectations and Monetary Policy Design: Evidence from the Laboratory." *Macroeconomic Dynamics* 22 (4): 1035–75.
- Sargent, Thomas J. 1993. *Bounded Rationality in Macroeconomics*. Oxford: Oxford University Press.
- Sargent, Thomas J. 1999. *The Conquest of American Inflation*. Princeton: Princeton University Press.
- Scheffer, Marten. 2009. *Critical Transitions in Nature and Society*. Princeton: Princeton University Press.
- Scheffer, Marten, et al. 2012. "Anticipating Critical Transitions." *Science* 338 (6105): 344–48.
- Schönhofer, Martin. 1999. "Chaotic Learning Equilibria." *Journal of Economic Theory* 89 (1): 1–20.
- Shiller, Robert J. 1981. "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71 (3): 421–36.
- Sims, Christopher A. 2010. "Rational Inattention and Monetary Economics." In *Handbook of Monetary Economics*, Vol. 3A, edited by Benjamin M. Friedman and Michael Woodford, 155–81. Amsterdam: North-Holland.
- Slobodyan, Sergey, and Raf Wouters. 2012a. "Learning in an Estimated Medium-Scale DSGE Model." *Journal of Economic Dynamics and Control* 36 (1): 26–46.
- Slobodyan, Sergey, and Raf Wouters. 2012b. "Learning

- in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models." *American Economic Journal: Macroeconomics* 4 (2): 65–101.
- Smets, Frank, and Raf Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1 (5): 1123–75.
- Smets, Frank, and Raf Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.
- Smith, Vernon L. 1962. "An Experimental Study of Competitive Market Behavior." *Journal of Political Economy* 70 (2): 111–37.
- Smith, Vernon L., Gerry L. Suchanek, and Arlington W. Williams. 1988. "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets." *Econometrica* 56 (5): 1119–51.
- Stiglitz, Joseph E. 2018. "Where Modern Macroeconomics Went Wrong." *Oxford Review Economic Policy* 34 (1–2): 70–106.
- Sutan, Angela, and Marc Willinger. 2009. "Guessing with Negative Feedback: An Experiment." *Journal of Economic Dynamics and Control* 33 (5): 1123–33.
- Trichet, Jean-Claude. 2010. "Reflections on the Nature of Monetary Policy, Nonstandard Measures, and Finance Theory." Opening address at the ECB Central Banking Conference Frankfurt, 18 November. <https://www.ecb.europa.eu/press/key/date/2010/html/sp101118.en.html>.
- Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185 (4157): 1124–31.
- Vines, David, and Samuel Wills. 2018. "The Rebuilding Macroeconomic Theory Project: An Analytical Assessment." *Oxford Review Economic Policy* 34 (1–2): 1–42.
- Wagner, Florian. 2014. "Expectations in Experiments." *Annual Review of Economics* 6: 421–43.
- West, Kenneth D. 1987. "A Specification Test for Speculative Bubbles." *Quarterly Journal of Economics* 102 (3): 553–80.
- Woodford, Michael. 1990. "Learning to Believe in Sunspots." *Econometrica* 58 (2): 277–307.
- Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.
- Woodford, Michael. 2010. "Robustly Optimal Monetary Policy with Near-Rational Expectations." *American Economic Review* 100 (1): 274–303.
- Woodford, Michael. 2013. "Macroeconomic Analysis without the Rational Expectations Hypothesis." *Annual Review of Economics* 5: 303–46.
- Woodford, Michael. 2019. "Monetary Policy Analysis When Planning Horizons Are Finite." In *NBER Macroeconomics Annual 2018*. Vol. 33, edited by Martin Eichenbaum and Jonathan A. Parker, 1–50. Chicago: University of Chicago Press.
- Zeeman, E. C. 1974. "On the Unstable Behaviour of Stock Exchanges." *Journal of Mathematical Economics* 1 (1): 39–49.