

Learning as a Rational Foundation for Macroeconomics and Finance*

George W. Evans

University of Oregon and University of Saint Andrews
and

Seppo Honkapohja
Bank of Finland

September 8, 2011

Abstract

Expectations play a central role in modern macroeconomics. The econometric learning approach, in line with the cognitive consistency principle, models agents as forming expectations by estimating and updating subjective forecasting models in real time. This approach provides a stability test for RE equilibria and a selection criterion in models with multiple equilibria. Further features of learning, such as discounting of older data, use of misspecified models, or heterogeneous choice by agents between competing models, generate novel learning dynamics. Empirical applications are reviewed and the roles of the planning horizon and structural knowledge are discussed. We develop several applications of learning to macroeconomic policy: the scope of Ricardian equivalence, appropriate specification of interest-rate rules, implementation of price-level targeting to achieve learning-stability of the optimal RE equilibrium and whether under learning price-level targeting can rule out the deflation trap at the zero-lower-bound.

*Any views expressed are those of the authors and do not necessarily reflect the views of the Bank of Finland. Financial support from National Science Foundation Grant no. SES-1025011 is gratefully acknowledged.

Key words: Cognitive consistency, E-stability, least-squares, persistent learning dynamics, business cycles, monetary policy, asset prices.

JEL classification: E32, D83, D84, C62

1 Introduction

Expectations play a central role in modern macroeconomics. Economic agents are assumed to be dynamic optimizers whose current economic decisions are the first stage of a dynamic plan. Thus households must be concerned with expected future incomes, employment, inflation, and taxes, as well as the expected trajectory of the stock market and the housing market. Firms must forecast the level of future product demand, wage costs, productivity levels, and foreign exchange rates. Monetary and fiscal policy-makers must forecast inflation and aggregate economic activity and consider both the direct impact of their policies and the indirect effect of policy rules on private-sector expectations.

The recent financial crisis has demonstrated that suddenly circumstances can change greatly and in such situations information and understanding become very imperfect. Even in normal times one can expect that agents are at best boundedly rational, and in exceptional times there will be a particular emphasis for agents to improve their knowledge of the situation, so that learning becomes central.

In this paper we discuss the central ideas about learning and bounded rationality using several standard macroeconomic models and settings (though we do not focus on finance and banking problems). One basic message is that in standard macroeconomic models rational expectations can emerge in the longer run, provided the agents' environment remains stationary for a sufficiently long period. However, from policy point of view, it is important to take into account the learning process and time periods when knowledge is quite imperfect and learning is a major driver of economic dynamics.

The structure of the paper is as follows. In the next section we develop the main methodological issues concerning expectation formation and learning, and discuss the circumstances in which rational expectations may arise. Section 3 reviews empirical work that applies learning to macroeconomic issues and asset pricing. In Section 4 we take up the implications of the use of structural information in learning and the form of the agents' decision rules. We then consider several applications: the scope of Ricardian equivalence,

appropriate specification of interest-rate rules, implementation of price-level targeting to achieve learning-stability of the optimal RE equilibrium and whether under learning commitment to price-level targeting can be sufficient to rule out the deflation trap zero interest-rate lower bound and return the economy to the intended rational expectations steady state.

2 Methodological issues in bounded rationality and learning

We develop the main ideas in an abstract setting. Macroeconomic models can often be summarized as a reduced-form multivariate dynamic system

$$y_t = \mathcal{F}(y_{t-1}, \{y_{t+j}^e\}_{j=0}^{\infty}, w_t, \eta_t), \quad (1)$$

where y_t is a vector of endogenous aggregate variables, and w_t is a vector of stochastic exogenous variables. Typically, w_t is assumed to follow a stationary stochastic process such as a finite-dimensional vector autoregression. The setting implicitly assumes that a representative-agent set-up is adequate: the vector y_t , which e.g. includes aggregate output, labor hours, consumption, inflation and factor prices, is the result of individual decisions, and aggregation to means is assumed acceptable. The literature contains a number of papers that allow for heterogeneity in agents' expectations or characteristics.

Crucially, y_t depends not only on the state of the system, captured by the exogenous variables and lagged endogenous variables, w_t and y_{t-1} , but also on expectations of future endogenous variables, $\{y_{t+j}^e\}_{j=1}^{\infty}$ and possibly on forecasts of the current endogenous variables. The presence of expectations $\{y_{t+j}^e\}_{j=0}^{\infty}$ is a key feature of the system that makes economics distinct from natural sciences.

At a general level, a learning mechanism is, for each period t , a mapping from the time t information set to the sequence of expectations of future (and possibly current) values of relevant variables together with an initial set of expectations of these variables. Some crucial aspects of the system to bear in mind are:

(i) The horizon for decisions and expectations. In some models only one-step ahead forecasts matter, while in others there is a long or infinite horizon.

(ii) Degree of structural information. Do agents know the whole structure, part of the structure or do they forecast using a reduced form? In the latter case do they know the correct functional form?

(iii) The precise information set on which expectations are based. Expectations may be based on all observables dated at time $t - 1$ or at time t . In the latter case w_t is assumed known at t and the current endogenous aggregate state y_t may or may not be known when agents make their decisions, so that y_t will depend also on “forecasts” of contemporaneous variables y_t^e . There may also be unobserved shocks.

(iv) A learning rule describing how expectations are formed over time. One can think of various ways for updating expectations and various standard statistical forecasting rules are special cases of this general formulation.

Since the work of Muth (1961), Lucas (1972), and Sargent (1973), the benchmark model of expectation formation in macroeconomics has been rational expectations. This posits, for both private agents and policy-makers, that expectations are equal to the true statistical conditional expectations of the unknown random variables. The “learning theory” approach in macroeconomics argues that although rational expectations is the natural benchmark, it is implausibly strong. We need a more realistic model of rationality, which may, however, be consistent with agents eventually learning to have rational expectations.

A natural criterion for a model of rationality is what we call the *cognitive consistency principle*, that economic agents should be assumed to be about as smart as (good) economists. This still leaves open various possibilities, since we could choose to model households and firms like economic theorists or, alternatively, model them like econometricians. The adaptive or econometric learning approach, which will here be our principal focus,¹ takes the latter viewpoint, arguing that economists, when they forecast future economic aggregates, usually do so using time-series econometric techniques. This seems particularly natural since neither private agents nor economists at central banks or other institutions know the true model. Instead economists formulate and estimate models. These models are re-estimated and possibly reformulated as new data become available. Economists themselves engage in processes of learning about the economy.

2.1 Least-Squares Learning and E-stability

We introduce the econometric learning approach using a simple linear model

$$y_t = \mu + \alpha y_t^e + \delta' w_{t-1} + \eta_t. \quad (2)$$

¹However, see also Section 4.2.1, where we discuss the eductive approach.

Here y_t is a scalar endogenous variable, w_{t-1} is a vector of exogenous observable variables and η_t is an unobservable random shock. The expectation variable $y_t^e = E_{t-1}^* y_t$ in this model is the expectation of y_t based on the set of observables dated $t - 1$ or earlier. The notation $E_{t-1}^* y_t$ here indicates the information set and the use of E_{t-1}^* instead of E_{t-1} indicates that the expectations of economic agents are not (necessarily) fully rational (in the usual sense of “rational expectations”). This model is particularly simple in that there are no expectations of future variables included. For convenience in this example we also date the exogenous observable variables w_{t-1} since the information set is taken to be variables at time $t - 1$.

The unique rational expectations equilibrium of this model is

$$y_t = \bar{a} + \bar{b}' w_{t-1} + \eta_t, \quad \bar{a} = (1 - \alpha)^{-1} \mu, \quad \bar{b} = (1 - \alpha)^{-1} \delta, \quad (3)$$

as is easily verified by applying the method of undetermined coefficients with the functional form $y_t = a + b' w_{t-1} + \eta_t$. Two well-known economic examples lead to reduced-form model (2).

Example 1: (Lucas-type aggregate supply model). A simple version of the “Lucas islands” model presented in Lucas (1973) consists of the aggregate supply and demand equations

$$\begin{aligned} q_t &= \bar{q} + \pi(p_t - p_t^e) + \zeta_t \\ m_t + v_t &= p_t + q_t, \end{aligned}$$

where $\pi > 0$. Here v_t is a velocity shock and m_t is the money supply. Assuming both v_t and the m_t rule depend in part on exogenous observables w_{t-1} , we have

$$\begin{aligned} v_t &= \mu + \gamma' w_{t-1} + \xi_t, \\ m_t &= \bar{m} + \rho' w_{t-1} + u_t, \end{aligned}$$

where u_t, ξ_t and ζ_t are white noise shocks. The reduced form of the model is of the form (2) with $y_t \equiv p_t$ and $0 < \alpha = \pi(1 + \pi)^{-1} < 1$ and $\eta_t = (1 + \pi)^{-1}(u_t + \xi_t - \zeta_t)$.

Example 2: (Muth market model). In the classic “cobweb model” analyzed under rational expectations by Muth (1961), an isolated market has a one-period production lag with competitive and (for simplicity) identical supply decisions based on expected price $p_t^e = E_{t-1}^* p_t$ the period before. Demand and supply are

$$d_t = m_I - m_p p_t + v_{1t}, \quad \text{and} \quad s_t = r_I + r_p p_t^e + r_w' w_{t-1} + v_{2t},$$

where v_{1t}, v_{2t} are white noise. With market clearing $d_t = s_t$, we obtain (2) as the reduced form with $y_t \equiv p_t$, $\eta_t = (v_{1t} - v_{2t})/m_p$, $\mu = (m_I - r_I)/m_p$, $\delta = -m_p^{-1}r_w$ and $\alpha = -r_p/m_p$. Note that $\alpha < 0$ for $r_p, m_p > 0$.

The rational expectations equilibrium (3) of the reduced form (2) implies the rational expectation

$$y_t^e = E_{t-1}y_t = \bar{a} + \bar{b}'w_{t-1}.$$

How would agents come to have rational expectations? In the econometric approach to learning, agents are assumed, like econometricians, to use past data to estimate the parameters a, b in the perceived model

$$y_t = a + b'w_{t-1} + \eta_t, \tag{4}$$

and to use the parameter estimates a_{t-1}, b_{t-1} to make forecasts

$$y_t^e = a_{t-1} + b_{t-1}'w_{t-1}.$$

For simplicity, all agents are here assumed to have the same expectations.

The exogenous shocks w_{t-1} and η_t and the expectations y_t^e determine actual y_t according to the model (2). This is called the *temporary equilibrium* at time t . Under learning the parameters of the forecasting model are updated in t to (a_t, b_t) , e.g. using least squares. The sequence of temporary equilibria under learning can thus be defined recursively. Agents are said to attain rational expectations asymptotically if $a_t, b_t \rightarrow \bar{a}, \bar{b}$ as $t \rightarrow \infty$.

For this set-up Bray and Savin (1986), Fourgeaud, Gourieroux, and Pradel (1986), and Marcet and Sargent (1989b) demonstrated asymptotic convergence to rational expectations with probability one if $\alpha < 1$ and convergence with probability zero if $\alpha > 1$. Why is the condition $\alpha < 1$ required for rational expectations to be attained? The set-up differs from the standard econometric assumptions in that the model is self-referential: because the actual price process y_t depends on y_t^e and because the coefficients of y_t^e evolve under learning, the system during the learning transition is actually non-stationary. Under least-squares learning agents neglect this nonstationarity and are thus making a subtle misspecification error. However, when $\alpha < 1$ the specification error vanishes asymptotically and the system converges to a stationary process: their estimates are econometrically consistent, and in fact converge with probability one. When instead $\alpha > 1$ the self-referential feature of the system dominates, agents' parameter estimates diverge and the rational expectations solution fails to be stable under learning.

Provided $\alpha < 1$, agents' expectations y_t^e converge to rational expectations asymptotically. Agents are only boundedly rational in that their forecasts have systematic forecast errors during the learning process, but we have a rational foundation for rational expectations in the sense that agents do not stick with incorrect parameters in the long run. By adjusting their estimated parameters in response to forecast errors using least squares, agents are eventually led to fully consistent expectations.

The econometric learning approach to expectation formation can thus be used to obtain an understanding of how agents might come to have rational expectations, and this thus provides a test of the plausibility of rational expectations in a particular model. The stability conditions are typically straightforward to obtain using the expectational stability (E-stability) technique, which examines the core dynamics of the differential equation approximation to the stochastic discrete-time system under learning. The E-stability technique looks at the actual law of motion generated by a given perceived law of motion. For the case at hand the actual law of motion is given by inserting the perceived law of motion expectations $y_t^e = a + b'w_{t-1}$ into (2), yielding the actual law of motion

$$y_t = \mu + \alpha a + (\delta + \alpha b)'w_{t-1} + \eta_t,$$

and a corresponding mapping from perceived law of motion to actual law of motion given by $T(a, b) = (\mu + \alpha a, \delta + \alpha b)$. E-stability is defined as stability of (\bar{a}, \bar{b}) under the ordinary differential equation $d(a, b)/d\tau = T(a, b) - (a, b)$, which in the current case immediately yields the correct condition $\alpha < 1$. Here τ denotes "virtual" time.²

The general E-stability technique is described in detail in Evans and Honkapohja (2001) and summarized in Evans and Honkapohja (2009b). Dynamic macroeconomic models, in which there are expectations of future endogenous variables, can have multiple equilibria, including in some cases sunspot equilibria or cycles, and E-stability can be used to select those rational expectations equilibrium that are attainable under econometric learning, e.g., see Evans, Honkapohja, and Romer (1998) and Evans, Honkapohja, and Marimon (2007). Thus, in principle the cognitive consistency principle and the bounded rationality econometric learning approach can provide a rational foundation for the rational expectations approach in macroeconomics,

²Virtual time can be related to calendar time measured in discrete periods. See, e.g. Evans and Honkapohja (2009b).

both in cases of “well-behaved” economies, with a unique equilibrium, and in some cases in which there are fully self-fulfilling, but inefficient, macroeconomic fluctuations.

However, the econometric learning approach can also generate new learning dynamics not found under the rational expectations hypothesis, and recent research has focused on a number of issues that are at the heart of the question “Are there rational foundations for macroeconomics?”

2.2 Structural change and misspecification

We argued above that the econometric learning approach could in some cases provide a rational foundation for the use of rational expectations in macroeconomics. But the approach also raises some natural questions that can undercut the rational expectations hypothesis. Real-world econometricians do not know the correct specification of the data generating process. Limitations on degrees of freedom imply that econometric forecasting models will typically be misspecified, e.g. underparameterized in terms of explanatory variables or lags. If this is true for applied econometricians then the cognitive consistency principle implies that it is also true for households, firms and policy-makers.

Furthermore, the economy continually undergoes structural change, whether incrementally or occasionally in sudden shifts, with the change taking an unknown form. If the structure of the economy does not remain stationary over time, the true data-generating may never be known. Appealing again to the cognitive consistency principle, economic agents should be aware of the likelihood of structural change and take measures to deal with it. Finally, if economic agents cannot be expected to know the correct specification of the data generating process, then agents are likely to disagree about the best forecasting model to use: heterogeneity of expectations can be expected to be widespread. We now outline approaches that incorporate these points.³

³In this volume Frydman and Goldberg (2010a) describe rational expectations models as “fully predetermined models” and also group the learning approach in the same camp. We feel this is misleading. While learning as a stability theory for RE was the main focus of the early literature on learning, there is an increasing focus on the dynamics introduced by learning. Learning models now focus on a wide variety of issues, many of which are described in this section. Topics like heterogeneity, misspecification and structural change show that the learning literature can account for aspects of “unknown unknowns” as well as “known unknowns”. Furthermore, the set of adaptive learning rules currently being

2.2.1 Misspecification and restricted perceptions equilibria

The reality that econometricians sometimes use misspecified models suggests we should consider implications of agents using misspecified econometric forecasting models. This can result in convergence to “restricted perceptions equilibria” in which the agents use the best misspecified model.

As a simple example, consider (2) and suppose the perceived law of motion takes the form of omitting a subset of the variables w_t . Specifically, write $w'_t = (w'_{1,t}, w'_{2,t})$ and assume that the agents’ perceived law of motion takes the form

$$p_t = a + c'w_{1,t-1} + \varepsilon_t,$$

where ε_t is believed to be white noise. This perceived law of motion gives an actual law of motion $p_t = (\mu + \alpha a) + (\gamma_1 + \alpha c)'w_{1,t-1} + \gamma_2'w_{2,t-1} + \eta_t$. For this actual law of motion the best model in the permitted class of perceived law of motions is the “projected actual law of motion” obtained by computing the linear projection of p_t onto the information set $E(p_t|1, w_{1,t-1}) = \tilde{T}_a(a, c) + \tilde{T}_c(a, c)'w_{1,t-1}$. This gives a mapping $(a, c) \rightarrow \tilde{T}(a, c)$, but now $\tilde{T}(a, c) = (\tilde{T}_a(a, c), \tilde{T}_c(a, c))$ depends on the covariance matrix for (w_{1t}, w_{2t}) . A fixed point of \tilde{T} has the property that forecasts are optimal relative to the restricted information set used by agents, and we therefore call this solution a *restricted perceptions equilibrium*. It can be shown that under least-squares learning there is convergence to the restricted perceptions equilibrium if $\alpha < 1$.⁴

2.2.2 Constant gain learning and escape dynamics

Suppose now that agents are concerned about structural change of an unknown form. A changing structure is tracked more effectively by weighting recent data more heavily. This can be conveniently done using a discounted (or “constant-gain”) version of least squares. Since constant-gain least squares weights recent data more heavily, convergence will be to a stochastic process near the rational expectations equilibrium, rather than to the rational expectations equilibrium itself. As an example, consider again the model (2). Under standard least squares each data point receives the same weight $1/t$.

explored is quite broad and thus allows for a wide range of possible learning dynamics.

⁴Restricted perceptions equilibrium was introduced in Chapter 13 of Evans and Honkapohja (2001). Related notions have been suggested by Marcat and Sargent (1989a), Sargent (1991), and Hommes and Sorger (1997).

The current data point thus has a declining weight as t increases. If agents instead use constant-gain least squares, the current data receive a fixed weight $0 < \gamma < 1$ and weights on past data points decline geometrically at rate $1 - \gamma$. For the model (2) it can be shown that estimates (a_t, b_t) now fail to converge fully to the rational expectations values. Instead the estimates converge to a stochastic process centered on (\bar{a}, \bar{b}) , with a finite variance scaled by the gain γ , so that learning remains imperfect.

In some cases this use of constant-gain least squares can have major implications for economic policy. For example, Cho, Williams, and Sargent (2002) show the possibility of escape dynamics in inflation models, in which parameter estimates for an extended period of time stray far from the rational expectations equilibrium. As another example, Orphanides and Williams (2007) argue that monetary policy needs to take account of imperfect learning by private agents.

2.2.3 Heterogeneous expectations

The preceding discussion has assumed homogeneous expectations for analytical convenience. In practice, heterogeneous expectations can be a major concern. In some models the presence of heterogeneous expectations does not have major effects on convergence conditions, as first suggested by Evans and Honkapohja (1996) and further studied by Evans, Honkapohja, and Marimon (2001) and Giannitsarou (2003). These papers assume that expectations are the only source of heterogeneity, i.e., agents' preferences and technologies are identical. Honkapohja and Mitra (2006) showed that interaction of structural and expectational heterogeneity can make the conditions for convergence of learning significantly more stringent than those obtained under homogeneous expectations. Agents' behavioral heterogeneity, speed of learning and the mixture of specific learning rules all affect the conditions for convergence to rational expectations equilibrium. There are also conditions on agents' characteristics that ensure convergence for all speeds of learning and mixtures of specific learning rules.

Heterogeneity of expectations can have other implications. For example, the papers just cited do not focus on multiple equilibria. Guse (2006) studies the stability of equilibria in a model with multiple solutions emphasizing the distribution of forecasting heterogeneity can play an important role in determining stability properties. Misspecification and heterogeneous expectations are combined in Berardi (2007) and Berardi (2008) to yield models of het-

erogeneous expectations equilibria. Heterogeneity of beliefs among agents is clearly central in empirical and experimental research on expectations formation, see Section 3.5 below for references.⁵

In many settings, there is no difficulty incorporating heterogeneous expectations into adaptive learning models, and this heterogeneity can take various forms. For example, Evans, Honkapohja, and Marimon (2001) allows for stochastic heterogeneity across agents' expectations due to random gains and random inertia (updating frequency). In (mean) constant gain versions of these rules this would lead to persistent heterogeneity. Expectation shocks can also be included, as in Evans and Honkapohja (2003a), Section 4, and Milani (2011), and it would be straightforward to allow for idiosyncratic components to these shocks.

2.2.4 Dynamic predictor selection

Once one accepts that forecasting models may be misspecified and that agents have heterogeneous expectations, one is driven toward the possibility that agents choose between competing models. If agents can alter their choices over time, then this gives rise to the “dynamic predictor selection” approach. Having multiple forecasting models in play is one way of obtaining heterogeneous expectations and can lead to a variety of learning dynamics, including regime switching behavior and additional volatility.

Brock and Hommes (1997) postulate that agents have a finite set of predictors or expectation functions, and that each predictor has a “fitness,” i.e. an estimate, based on past data, of its profits net of the cost of using the predictor. The proportion of agents who select a predictor depends on this fitness. Brock and Hommes (1997) study in detail the resulting “adaptively rational expectations dynamics” for the standard cobweb model when there are two predictors: a costly rational predictor and a costless naive forecast. They show that cycles and even chaotic dynamics can arise in such a setting.

The dynamic predictor selector framework is extended by Branch and Evans (2006a), Branch and Evans (2007) and Branch and Evans (2010) to incorporate stochastic features and econometric learning. Noting pervasive degrees of freedom limitations, Branch and Evans appeal to the merits of par-

⁵An important approach to heterogeneous expectations, developed by M. Kurz in a number of papers, is the concept of rational beliefs equilibrium, which requires consistency of heterogeneous beliefs with the empirical distribution of past data. See for example Kurz (1997), Kurz (2009), and Kurz (2011).

simonious forecasting models and study the implications of agents choosing between (equally costly) underparameterized models. As a simple illustration, each of two competing models might omit one of the two exogenous shocks. In models with negative expectational feedback, e.g. model (2) with $\alpha < 0$ as in the cobweb model, there is the possibility of “intrinsic heterogeneity” in which both forecasting models are in use in equilibrium. In models with positive feedback, e.g. (2) with $\alpha > 0$ as in the Lucas-type monetary model, two “misspecification equilibria” can exist and agents may coordinate on either of the two forecasting models. Under real-time learning the stochastic process for inflation and output can then exhibit regime-switching or parameter drift, in line with much macroeconomic evidence.

The dynamic predictor selection approach has been used in other applications. For example, Brazier, Harrison, King, and Yates (2008) and De Grauwe (2011), both of which emphasize switching between alternative forecasting models, look at the implications of regime switching for macroeconomic policy. Another approach to model specification is for agents to switch between model specifications over time based on econometric criteria, as in Markiewicz (2010). A related approach, based on tests for misspecification, is developed in Cho and Kasa (2010).

Another natural approach when multiple models are in play would be for individual agents to do some form of averaging across forecasting models. A simple alternative would be to assume that agents have fixed weights between models as in the “natural expectations” employed by Fuster, Laibson, and Mendel (2010), in which agents use a fixed weight to average between rational expectations and an “intuitive model”.⁶ A contrasting formulation is to assume that agents update model estimates using standard econometric tools and then use Bayesian model averaging, e.g. as in Slobodyan and Wouters (2008).

The central message of the modified learning procedures, discussed in this section, is that variations of econometric learning, which adhere to the cognitive consistency principle and reflect real-world concerns of applied econometricians, can lead to new learning dynamics that are qualitatively different from rational expectations dynamics. In these learning dynamics agents are boundedly rational, in that their perceived law of motion does not *fully* reflect the actual economic dynamics, but given their knowledge, they are

⁶From the viewpoint of cognitive consistency both the weight between and the parameters of the forecasting models could be allowed to respond to data over time.

forecasting in a nearly optimal way.

3 Learning and Empirical Research

There is an expanding literature that employs the learning approach to study empirical issues in macroeconomics and finance. We give an overview of the main topics that have been studied.

3.1 Rise and fall of inflation

Several recent papers have argued that learning plays a central role in the historical explanation of the rise and fall of US inflation over the 1960-1990 period. Sargent (1999) and Cho, Williams, and Sargent (2002) emphasize the role of policy-maker learning. They argue that if monetary policy-makers attempt to implement optimal policy while estimating and updating the coefficients of a misspecified Phillips curve, there will be both periods of inefficiently high inflation and occasional escapes to low inflation. Sargent, Williams, and Zha (2006) estimate a version of this model. They find that shocks in the 1970s led the monetary authority to perceive a trade-off between inflation and unemployment, leading to high inflation, and subsequent changed beliefs about this trade-off account for the conquest of US inflation during the Volcker period.

Primiceri (2006) makes a related argument, emphasizing both policy-maker learning about the Phillips curve parameters and the aggregate demand relationship, and uncertainty about the unobserved natural rate of unemployment. The great inflation of 1970s initially resulted from a combination of underestimates of both the persistence of inflation and the natural rate of unemployment. This also led policy-makers to underestimate the disinflationary impact of unemployment.

Other empirical accounts of the period that emphasize learning include Bullard and Eusepi (2005), which examines the implications of policy-maker learning about the growth rate of potential output, Orphanides and Williams (2005a), which underscores both private-agent learning and policy-maker misestimates of the natural rate of unemployment, Orphanides and Williams (2005c), which looks at estimated models that focus on the explanation of the large increase in inflation rates in the 1970s, and Cogley and Sargent (2005), which develops a historical account of inflation policy emphasizing

Bayesian model averaging and learning by policy-makers uncertain about the true economic model.

Recent papers include Ellison and Yates (2007) and Carboni and Ellison (2008), which emphasize the importance of policy-maker model uncertainty and the role of central bank learning in explaining the historical evolution of inflation and unemployment in the post 1950 period.

3.2 Latin American inflation

Marcet and Nicolini (2003) use an open-economy extension of the standard seigniorage model of inflation, in which government spending is financed by printing money. They present a calibrated learning model that aims to explain the central stylized facts about hyperinflation episodes during the 1980s in a number of South American countries: (i) recurrence of hyperinflation episodes, (ii) exchange rate rules stop hyperinflations, though new hyperinflations eventually occur, (iii) during a hyperinflation, seigniorage and inflation are not highly correlated, and (iv) average inflation and seigniorage are strongly positively correlated across countries.

These facts are difficult to explain using the rational expectations assumption. Under learning there are occasional escapes from the low inflation steady state to an unstable hyperinflationary process that is eventually arrested by imposing an exchange rate rule. All four stylized facts listed above can be matched using this model. For example, under learning higher levels of deficits financed by seigniorage make average inflation higher and the frequency of hyperinflations greater. Simulations of a calibrated model look very plausible.

3.3 Real business cycle applications

Williams (2004) explores the real business cycle model dynamics under learning. Using simulations he shows that the dynamics under rational expectations and learning are not very different unless agents need to estimate structural aspects as well as the reduced form perceived law of motion parameters. Huang, Liu, and Zha (2009) focus on the role of misspecified beliefs and suggest that these can substantially amplify the fluctuations due to technology shocks in the standard real business cycle model. Eusepi and Preston (2011) incorporate infinite-horizon learning into the standard real business cycle

model and find that under learning the volatilities and persistence of output and employment are higher under learning than under rational expectations.

Other papers on learning and business cycle dynamics include Van Nieuwerburgh and Veldkamp (2006) and Giannitsarou (2006). The former formulates a model of Bayesian learning about productivity and suggests that the resulting model can explain the sharp downturns that are an empirical characteristic of business cycles. The latter extends the basic real business cycle model to include government spending financed by capital and labour taxes. It is shown that if a reduction of capital taxes is introduced following negative productivity shocks, the learning adjustment exhibits a delayed response in economic activity, in contrast to an immediate positive response under rational expectations.

3.4 Asset pricing and learning

The initial work by Timmermann (1993), Timmermann (1996) studied the implications of learning in the standard risk-neutral asset pricing model. The main finding was that in both short- and long-horizon models learning increased the volatility of asset prices during learning.

In more recent literature on learning and stock prices Brock and Hommes (1998) introduce heterogeneous expectations using the dynamic predictor selection methodology discussed earlier. Branch and Evans (2011), Adam, Marcet, and Nicolini (2008) and Lansing (2010) present models in which stock market bubbles arise endogenously.

Branch and Evans (2011) examine learning within a portfolio model in which the demand for risky asset depends positively on the expected returns and negatively on expected conditional variance of returns. Under constant-gain learning there is a regime in which stock prices exhibit bubbles and crashes driven by changing estimates of risk. Adam, Marcet, and Nicolini (2008) use the standard consumption-based model of stock prices modified to include learning. The model exhibits a number of salient features of the data including mean reversion of returns, excess volatility, and persistence of price-dividend ratios. A calibrated version of the model is shown to match many aspects of US data. Lansing (2010) shows that in Lucas-type asset pricing models, there are driftless near-rational solutions that are stable under learning, and which generate intermittent bubbles and dynamics qualitatively similar to long-run US stock market data.

The model of LeBaron (2010) focuses on the role of heterogeneous gains

in learning rules for estimating the mean returns and conditional variances by risk averse investors. He argues that agents putting relatively large weights on recent past are important for volatility magnification of asset price volatility and replication of long samples of U.S. financial data.

In summary, several recent papers are arguing that adaptive learning can play a key role in explaining asset-price behavior. The issues of market completeness and the role of financial derivatives has also received some attention. Using the eductive approach Guesnerie and Rochet (1993) demonstrate that opening futures markets can be destabilizing, and using a dynamic predictor selection approach, Brock, Hommes, and Wagener (2009) show that adding hedging instruments can destabilize markets and increase price volatility.

Exchange rate dynamics also exhibit a number of puzzles that learning models may resolve. For example, Kim (2009) shows that adaptive learning can generate the excess volatility, long swings and persistence that appears the data. Chakraborty and Evans (2008) focus on the forward-premium puzzle, and argue that adaptive learning can explain this anomaly while simultaneously replicating other features of the data such as positive serial correlation of the forward premium. Another, potentially complementary, approach to exchange-rate modeling is based on dynamic predictor selection, see De Grauwe and Grimaldi (2006). Further applications of learning to exchange rates include Kasa (2004), Mark (2007), and Markiewicz (2010).⁷

3.5 Estimated models with learning

The rational expectations version of New Keynesian models need to incorporate various sources of inertia arising from indexation and habit-formation in preferences to account for observed persistence in inflation and output

⁷In a number of publications Roman Frydman and Michael Goldberg have recently developed “imperfect knowledge economics” as a model of “non-routine change” in expectations formation. Frydman and Goldberg (2010b) describe the application of this approach to account for persistence and long swings in asset prices and exchange rates. It is difficult to assess their approach (presented in their section 6.2) since several key aspects are purposely only loosely specified. These include updating of forecasting strategy $\Delta\beta_t^i$, the bound δ_t for change in baseline drift, and some aspects of the stochastic structure. Juselius (2010) finds that many empirical regularities in exchange rate dynamics under rational expectations are empirically violated and suggests that the Frydman-Golberg approach works better. However, she does not relate her empirical regularities to those implied by other learning and related models that relax the rational expectations assumption.

data. These have been criticized as being *ad hoc*. Incorporating learning provides an alternative to account for the observed persistence. This point was initially made in a simple calibrated model by Orphanides and Williams (2005b). Milani (2007) addresses this issue using an estimated DSGE model with learning and finds that both habit formation and indexation inertia have minor roles when estimation allows for adaptive learning. The implications of incorporating learning within applied DSGE models has most recently been explored by Slobodyan and Wouters (2007) and Slobodyan and Wouters (2008).

There is also empirical work on forecasts, based on survey data, and indirect measures of expectations from asset markets that help to assess the alternative models of learning and expectations formation. For recent papers, see Branch (2004), Branch and Evans (2006b), Orphanides and Williams (2005a), Basdevant (2005), Pfajfar (2007), and Pfajfar and Santoro (2007). For research on expectations formation and learning in experiment settings, see Marimon and Sunder (1993), Marimon and Sunder (1995), Evans, Honkapohja, and Marimon (2001), Adam (2007), and the new survey Hommes (2011).

Recently, Milani (2011) has investigated the importance of expectations as a driving force for business cycles in the United States. In an estimated New Keynesian model with constant-gain VAR(1) learning, Milani (2011) uses survey data on expectations in conjunction with aggregate macro data both to estimate the structural parameters of the model and to identify expectation shocks, interpreted as arising from shifts in market sentiment. His provocative conclusion is that expectations shocks “can account for roughly half of business cycle fluctuations.”

We think that empirically oriented research on learning will continue to grow. As discussed above, adaptive learning has the potential to resolve a number of puzzles and difficulties that rational expectations models encounter when confronted with the data.

4 Further Issues in Modeling Learning

4.1 The planning horizon

In the Lucas/Muth model and in overlapping generations models with two-period lifetimes, agents in the current period make forecasts for the values

of aggregate variables in the next period. However, many modern macroeconomic models are set in a representative-agent framework with infinitely-lived agents who solve infinite-horizon dynamic optimization problems. Typically, under rational expectations the reduced-form equations for these models can be stated in the form $y_t = \mathcal{F}(y_{t-1}, y_{t+1}^e, w_t, \eta_t)$, where $y_{t+1}^e = E_t y_{t+1}$. This reduction relies on the use of the Euler equations to describe the first-order optimality conditions.

There are alternative approaches to learning and bounded rationality in infinite-horizon settings. In Evans and Honkapohja (2001), Chapter 10, the learning framework was kept close to the rational expectations reduced-form set-up, a procedure that can be justified if agents make decisions based directly on their Euler equations. This approach has been used, for example, in Bullard and Mitra (2002) and Evans and Honkapohja (2003b). An alternative approach, followed by Preston (2005), Preston (2006), assumes that households use estimated models to forecast aggregate quantities infinitely far into the future to solve for their current decisions.⁸ We now illustrate the two approaches using a simple endowment economy.⁹

A representative consumer makes consumption-saving decisions using the intertemporal utility function

$$E_t^* \sum_{s=t}^{\infty} \beta^{s-t} U(C_s). \quad (5)$$

Each agent has a random endowment Y_t of the perishable good. There is a market in safe one-period loans with gross rate of return R_t , known at t . Initial wealth for each agent is zero. Output Y_s follows an exogenous process given by

$$\log Y_s = \mu + \rho \log Y_{s-1} + v_s,$$

where $|\rho| < 1$ and v_s is white noise. Expectations are not necessarily rational, indicated by $*$ in the expectations operator. Defining $\mathcal{R}_{t+1,s}^{-1} = \prod_{j=t+1}^s R_{t+j}$, the household's intertemporal budget constraint is

$$C_t + \sum_{s=t+1}^{\infty} \mathcal{R}_{t+1,s} C_s = Y_t + \sum_{s=t+1}^{\infty} \mathcal{R}_{t+1,s} Y_s.$$

⁸Infinite-horizon learning based on an iterated Euler equation was used by Sargent (1993), pp. 122-125, in the “investment under uncertainty” model. See also example e of Marcet and Sargent (1989b).

⁹The passage is based largely on Honkapohja, Mitra, and Evans (2002), and we also draw on Evans, Honkapohja, and Mitra (2009).

Maximizing (5) subject to the intertemporal budget constraint yields the Euler necessary first-order condition, $U'(C_t) = \beta R_t E_t^* U'(C_{t+1})$. In the “Euler-equation learning” approach, the Euler equation is treated as a behavioral equation, determining for each agent their temporary equilibrium demand for C_t as a function of R_t and their forecast $E_t^* U'(C_{t+1})$. Imposing the market clearing condition $C_t = Y_t$, and using the representative agent setting, we obtain the temporary equilibrium interest rate $R_t^{-1} = \beta(E_t^* U'(C_{t+1}))/U'(Y_t)$.

Writing $y_t = \log(Y_t/\bar{Y})$, etc. and using log-linearizations around $\bar{Y} = \bar{C} = (1 - \rho)^{-1}\mu$, $\bar{R} = \beta^{-1}$ yields $y_t = \rho y_{t-1} + v_t$ and the consumption demand

$$c_t = E_t^* c_{t+1} - \sigma r_t, \quad (6)$$

where $\sigma = -U'(\bar{C})/(U''(\bar{C})\bar{C})$. In the temporary equilibrium $c_t = y_t$ and $r_t = \sigma^{-1}(E_t^* c_{t+1} - y_t)$. The rational expectations equilibrium of the linearized model is given by

$$r_t = -(1 - \rho)\sigma^{-1}y_t \text{ and } E_t c_{t+1} = \rho y_t.$$

To formulate “Euler-equation” learning, based on (6), suppose agents have the perceived law of motion

$$E_t^* c_{t+1} = m + n y_t, \quad (7)$$

with coefficient estimates (m_t, n_t) obtained using a regression of c_s on y_{s-1} using data $s = 1, \dots, t-1$. As usual, (m_t, n_t) are updated over time. Will agents learn to hold rational expectations over time, i.e. will we have $(m_t, n_t) \rightarrow (0, \rho)$ as $t \rightarrow \infty$? This can be assessed using E-stability: the perceived law of motion (7) leads to the actual law of motion $r_t = -\sigma^{-1}[y_t(1 - n) - m]$ and $c_t = y_t$. Since the actual law of motion forecasts are $E_t c_{t+1} = \rho y_t$, the T-map is simply $T(m, n) = (0, \rho)$. The E-stability differential equation $d(m, n)/d\tau = T(m, n) - (m, n)$ is stable, and there is convergence of least-squares learning to rational expectations in this model.

To summarize, under Euler-equation learning, agents choose their c_t using (6). This requires a forecast of the agent’s own c_{t+1} . This forecast is made via (7), in which agents assume that their future consumption is related (as in the rational expectations equilibrium) to the key state variable y_t . Thinking one step ahead in this way appears to us to be a plausible and natural form of bounded rationality. Although this formulation does not

explicitly impose the intertemporal budget constraint, it can be verified that along the learning path both the intertemporal budget constraint and the transversality condition are satisfied.¹⁰

An alternative approach treats consumption demand each period as based on forecasts over an infinite horizon.¹¹ We call this approach, presented for the New Keynesian model in Preston (2005), infinite-horizon learning, and we describe it for the current context. Log-linearizing the intertemporal budget constraint yields $c_t + \sum_{s=t+1}^{\infty} \beta^{s-t} E_t^* c_s = y_t + \sum_{s=t+1}^{\infty} \beta^{s-t} E_t^* y_s$. From the linearized Euler equation (6) we have $E_t^* c_s = c_t + \sigma \sum_{j=t}^{s-1} E_t^* r_j$ for $s \geq t+1$. Substituting into the linearized intertemporal budget constraint and solving for c_t leads to the behavioral equation

$$c_t = (1 - \beta)y_t - \sigma\beta r_t + \sum_{s=t+1}^{\infty} \beta^{s-t} [(1 - \beta)E_t^* y_s - \sigma\beta E_t^* r_s], \quad (8)$$

where we have assumed that both y_t and r_t are known at t .

Under infinite-horizon learning, suppose agents do not know the rational expectations relationship between y_t and r_t , but have the perceived law of motion

$$r_t = d + f y_t,$$

where at time t the coefficients are estimated to be d_t, f_t . To determine E-stability, use $E_t^* r_s = d + f y_s$ and $E_t^* y_s = \rho^{s-t} y_t$ in (8), impose market clearing $c_t = y_t$, and solve for r_t to obtain the implied actual law of motion, given by

$$T(d, f) = (-\beta(1 - \beta)^{-1}d, -(1 - \beta\rho)^{-1}(\sigma^{-1}(1 - \rho) + \beta\rho f)).$$

The fixed point of T is the rational expectations equilibrium and it is easily checked that the E-stability differential equation is stable. Thus, we have convergence of least-squares learning to rational expectations under infinite-horizon learning.

Although for this particular model, learning stability holds for both Euler-equation and infinite-horizon learning, in more general models it is possible for stability to depend on the planning horizon of the agents. For example, in

¹⁰Euler-equation learning is a special case of shadow price learning, which can be shown to deliver asymptotically optimal decision-making in general settings. See Evans and McGough (2010).

¹¹Thus, infinite-horizon agents explicitly solve dynamic optimization problems, which can be viewed as a version of the “anticipated utility” approach formulated by Kreps (1998) and discussed in Sargent (1999), and Cogley and Sargent (2008).

an real business cycle model with a mixture of rational and boundedly rational agents, Branch and McGough (2011) show that hump-shaped responses of consumption to productivity shocks can arise and are particularly strong when the boundedly rational agents have a long, finite planning horizon.¹² As another example, Eusepi and Preston (2011) show the implications for macroeconomic volatility of constant-gain infinite-horizon learning in a real business cycle setting.

4.2 Structural knowledge

In Section 2.1 we assumed that agents estimated the correct perceived law of motion, i.e. the form of the perceived law of motion corresponding to the rational expectations solution. Implicitly we assumed that they knew the list of observable variables that drive the rational expectations equilibrium, and that the solution was linear. However, we did not assume that they had the structural knowledge needed to compute the rational expectations equilibrium. Under (standard, i.e. decreasing gain) least-squares learning their expectations can converge to rational expectations (if $\alpha < 1$), even though they do not necessarily know the full economic structure. This is possible asymptotically because in order to forecast optimally all that is required in this setting is the linear projection of y_t onto the information set, and this can be consistently estimated by least squares. Of course, since the result is asymptotic, agents will not have rational expectations during the learning transition. This result raises the following question: if agents do have structural knowledge, will they be able to coordinate on rational expectations more quickly? Alternatively, if agents have incomplete structural knowledge, is there a natural way for them to incorporate this knowledge into econometric learning? We take up these two issues in turn.

4.2.1 Eductive stability under full structural knowledge

Consider again the reduced form model (2). For convenience we now omit the observable stochastic shocks, so that we have

$$y_t = \mu + \alpha y_t^e + \eta_t, \tag{9}$$

¹²Learning with finite planning horizons are developed further in Branch, Evans, and McGough (2010) and ?.

where η_t is a white noise exogenous shock. We now suppose that all agents know the structure (9) and that this is “common knowledge.” We further suppose that all agents are fully rational and know that all other agents are fully rational. Our cognitive consistency principle thus now takes a different form, in which we model economic agents like economic theorists. This leads to the “eductive” learning approach.

We now take up the eductive viewpoint for the model (2). The argument here was initially given by Guesnerie (1992) in the context of the cobweb model. See Evans and Guesnerie (1993) for the multivariate formulation and Guesnerie (2002) for a more general discussion using the eductive approach.

To focus the discussion, we use the cobweb example, which can be reformulated as a producers’ game in which the strategy of each firm is its output and the optimal choice of output depends on expected price. We assume that firms have identical costs. We allow for heterogeneous expectations, however, so that the equilibrium market price is given by

$$y_t = \mu + \alpha \int E_{t-1}^* y_t(\omega) d\omega + \eta_t,$$

where we now assume a continuum of agents indexed by ω and $E_{t-1}^* y_t(\omega)$ is the expectation of the market price held by agent ω . The rational expectations equilibrium is $y_t = \bar{a} + \eta_t$, where $\bar{a} = (1 - \alpha)^{-1}\mu$, and expectations are $E_{t-1} y_t = \bar{a}$. We now ask whether rational agents would necessarily coordinate on rational expectations.

The eductive argument works as follows. Let $S(\bar{a})$ denote a neighborhood of \bar{a} . Suppose it is common knowledge that $E_{t-1}^* p_t(\omega) \in S(\bar{a})$ for all ω . Then it follows that it is common knowledge that $E p_t \in |\alpha| S(\bar{a})$. Hence, by individual rationality, it is common knowledge that $E_{t-1}^* p_t(\omega) \in |\alpha| S(\bar{a})$ for all ω . If $|\alpha| < 1$ then this reinforces and tightens the common knowledge. Iterating this argument it follows that $E_{t-1}^* p_t(\omega) \in |\alpha|^N S(\bar{a})$ for all $N = 0, 1, 2, \dots$, and hence the rational expectations equilibrium $E p_t = \bar{a}$ is itself common knowledge. Guesnerie calls such a rational expectations equilibrium “strongly rational.” We also use the equivalent terminology that the rational expectations equilibrium is “eductively stable” or “stable under eductive learning.” We thus have the result: If $|\alpha| < 1$ then the rational expectations equilibrium is stable under eductive learning, while if $|\alpha| > 1$ the rational expectations equilibrium is not eductively stable.

Note two crucial differences from the least-squares adaptive learning results. First, the learning here takes place in mental time, not real time.

Given the common knowledge assumptions and full power of reasoning ability, if $|\alpha| < 1$ then rational agents would coordinate instantaneously, through a process of reasoning, on the rational expectations equilibrium. For the Lucas supply model this condition is always satisfied and for the cobweb model, with $\alpha < 0$, satisfaction of the stability condition depends on the relative slopes of the supply and demand curves and is satisfied when $\alpha > -1$. Second, when $\alpha < -1$ the rational expectations equilibrium is not eductively stable, but is asymptotically stable under adaptive learning.

The finding that eductive stability can be more demanding than stability under adaptive learning appears to be general. In simple models the eductive stability condition reduces to iterative E-stability, i.e. stability of the rational expectations equilibrium under iterations of the T-map, which itself is more demanding than E-stability. With structural heterogeneity and in dynamic models, the eductive stability conditions are even tighter, e.g. see Evans and Guesnerie (2003). A particularly striking example of this is the generic failure of strong eductive stability in infinite-horizon real business cycle models, established in Evans, Guesnerie, and McGough (2010).

4.2.2 Partial structural knowledge

In practical policy situations a question that often arises concerns the impact of anticipated future changes in policy. For instance, it is well recognized that there are long lags involved in changing fiscal policy. The process of changing taxes involves legislative lags, between when the new tax is proposed and when it is passed, and implementation lags, between when the legislation is signed into law and when it actually takes effect. The bulk of the literature on adaptive learning has focused on situations where the environment is stationary, so that in particular the policy used by the authorities never changes. Some papers have studied the effect of policy changes, but under the assumption that the policy change is completely *unanticipated* and agents begin to learn the new equilibrium as data arrive after the policy change. Such changes are examined in Evans, Honkapohja, and Marimon (2001), Marcet and Nicolini (2003) and Giannitsarou (2006). However, the anticipation of policy changes is likely to influence economic decisions even before the actual implementation of the proposed policy change.¹³

¹³One of the contributions of the rational expectations revolution was the idea that agents look forward and can anticipate the effects of an announced future shift in policy. Early examples are Sargent and Wallace (1973b) and Sargent and Wallace (1973a). The

Evans, Honkapohja, and Mitra (2009) propose a learning model in which agents combine limited structural knowledge about the economy with adaptive learning for other variables that they need to forecast. On this approach agents use statistical knowledge to forecast many economic variables, e.g. GDP growth and inflation, while incorporating structural knowledge about specific variables, e.g. announced future changes in government spending or taxes.¹⁴ In this setting, anticipated policy changes lead to immediate changes in the behavior of agents who are learning adaptively, even before the implementation of the proposed policy. Evans, Honkapohja, and Mitra (2009) show that the dynamic paths that result from this framework can differ significantly from the corresponding rational expectations path. The assumption that private agents know the monetary policy rule in models by Eusepi and Preston (2007) and Evans and Honkapohja (2010) are two other examples in which agents have partial structural knowledge but need to learn about other aspects of the economy.

Application: Ricardian Equivalence when expectations are not rational One of the most prominent theories in macroeconomics is the Ricardian Equivalence proposition that if taxes are non-distortionary then the mix of tax and debt financing of government purchases have no impact on the equilibrium sequence of key real variables. Conditions for validity or failure of the Ricardian proposition have been examined in the voluminous theoretical and empirical literature, e.g., see the survey papers by Bernheim (1987), Barro (1989), Seater (1993), and Ricciuti (2003).

A key assumption that has not been examined in this literature is the role of rational expectations. If expectations are made using adaptive (or statistical) learning rules, can Ricardian Equivalence still hold? Recently, Evans, Honkapohja, and Mitra (2010) have argued that Ricardian Equivalence holds under the usual conditions when agents are dynamic optimizers but with non-rational forecasts. Two key assumption are that agents understand the government’s budget constraint and expectations are based on a suitable information set.

The main results of Evans, Honkapohja, and Mitra (2010) are obtained in

rational expectations analysis of anticipated impacts is nowadays standard in textbooks.

¹⁴An alternative, and potentially complementary, approach is the “active cognition” framework of Evans and Ramey (1992) and Evans and Ramey (1998), in which agents employ a calculation algorithm based on a structural model, but are impeded by calculation costs.

the context of a standard Ramsey model with government bonds and lump-sum taxes. Ricardian Equivalence is often analyzed using this framework. The model is assumed to be non-stochastic and populated a large number of identical households (but individual households do not know that they are identical). Taxes are assumed to be lump-sum. At each time t the household maximizes their utility subject to a flow budget constraint and to No Ponzi Game and transversality conditions. The model implies a consumption function for households, which depends on current asset values and present values of labor incomes and taxes.¹⁵

The model also has a standard production function with labor and capital as inputs. The government's flow budget constraint states that end-of-period debt equals current gross interest payments on beginning-of-period debt plus the difference between government spending and tax receipts. The model also includes a standard market clearing equation.

Given pre-determined variables, current fiscal policy variables and expectations, a temporary equilibrium at time t is defined by the consumption function, the wage rate, the interest rate, the government flow budget constraint, and market clearing. Two key assumptions concerning households' perceptions of the government budget constraint are: (i) households understand the flow budget constraint of the government, and (ii) they believe that the expected limiting present value of government debt is zero.

These assumptions imply that the consumption function can be written as a function of the sum of current (gross) income from capital and the present values of wages and government spending. It follows that Ricardian Equivalence holds in the temporary equilibrium under the additional assumption that neither government spending nor expectations depend on current government financing variables.

The evolution of the economy over time is described as a sequence of temporary equilibrium with learning. The economy starts with some initial capital stock, public debt and beliefs about the future path of the economy and evolves along a path of temporary equilibria, given fiscal policy rules that determine government spending and taxes as well as debt dynamics. To close the dynamic model, a learning mechanism and its information set must be specified. The former is a mapping from the time t information set to the sequence of expectations over the infinite future, together with an initial set

¹⁵The present values are calculated using expected interest rates. All relevant expected present value sums are assumed to be finite.

of expectations. The latter is assumed to consist of observable variables and past expectations. The key result of Evans, Honkapohja, and Mitra (2010) is:

Proposition 1 *Assume that neither government spending nor expectations depend on current government financing variables (taxes and end-of-period debt). The Ramsey model exhibits Ricardian Equivalence, i.e., for all initial conditions, the sequence of consumption, capital, rates of return and wages along the path of equilibria with learning is independent of the government financing policy.*

For the result, it is crucial that the expectations of agents do not depend on government financing variables in addition to the usual assumptions about government spending and taxes. Evans, Honkapohja, and Mitra (2010) gives simple examples illustrating the role played by the assumption about expectations when agents are learning.

5 Learning and Monetary Policy

In the analysis of economic policy the rational expectations hypothesis should not be taken for granted, since expectations can be out of equilibrium at least for a period of time. Economic policies should in part be designed to avoid instabilities that can arise from expectational errors and the corrective behavior of economic agents in the face of such errors. We now consider aspects of these concerns for analysis of monetary policy in the widely-used New Keynesian framework.¹⁶

We first look at the implications of requiring stability under learning for the choice of the optimal interest-rate rule in the linearized New Keynesian model. We also consider price-level targeting from the learning viewpoint by deriving the optimal learnable policy rule. Second, we consider the performance of a Wicksellian price-level targeting rule in a global setting in which the zero lower bound can impose a constraint on interest-rate setting.

¹⁶For surveys of the growing literature on learning and monetary policy see Evans and Honkapohja (2003a), Bullard (2006), and Evans and Honkapohja (2009a).

5.1 Learning and the Choice of the Interest Rate Rule

5.1.1 The Basic Model

We use a linearized New Keynesian model that is very commonly employed in the literature, see Clarida, Gali, and Gertler (1999) for this particular formulation and references to the literature. The original nonlinear framework is based on a representative consumer, a continuum of firms producing differentiated goods under monopolistic competition and price stickiness.

The behavior of the private sector is described by two equations

$$x_t = -\varphi(i_t - E_t^* \pi_{t+1}) + E_t^* x_{t+1} + g_t, \quad (10)$$

which is the “IS” curve derived from the Euler equation for consumer optimization, and

$$\pi_t = \lambda x_t + \beta E_t^* \pi_{t+1} + u_t, \quad (11)$$

which is the price-setting rule for the monopolistically competitive firms.

Here x_t and π_t denote the output gap and inflation for period t , respectively. i_t is the nominal interest rate.¹⁷ $E_t^* x_{t+1}$ and $E_t^* \pi_{t+1}$ denote the private sector expectations of the output gap and inflation next period. These expectations need not be rational (E_t without * denotes rational expectations). The parameters φ and λ are positive and β is the discount factor so that $0 < \beta < 1$.

The shocks g_t and u_t are assumed to be observable and follow

$$\begin{pmatrix} g_t \\ u_t \end{pmatrix} = F \begin{pmatrix} g_{t-1} \\ u_{t-1} \end{pmatrix} + \begin{pmatrix} \tilde{g}_t \\ \tilde{u}_t \end{pmatrix}, \text{ where } F = \begin{pmatrix} \mu & 0 \\ 0 & \rho \end{pmatrix}. \quad (12)$$

$0 < |\mu| < 1$, $0 < |\rho| < 1$ and $\tilde{g}_t \sim iid(0, \sigma_g^2)$, $\tilde{u}_t \sim iid(0, \sigma_u^2)$ are independent white noise. The u_t shock is important for the policy issues since the g_t shock can be fully offset by appropriate interest-rate setting. For simplicity, μ and ρ are assumed to be known (if not, they could be estimated).

One part of the literature focuses on simple policy rules. Under Euler-equation learning, $E_t^* \pi_{t+1}$ and $E_t^* x_{t+1}$ represent private sector forecasts, which need not be rational, and (10)-(11) are interpreted as behavioral equations resulting from Euler-equation based decision rules. This set-up has been studied by Bullard and Mitra (2002) and Evans and Honkapohja (2003b) for

¹⁷Variables are expressed as deviations from their nonstochastic steady state values.

learning stability of the rational expectations equilibrium under alternative interest-rate rules.

Bullard and Mitra (2002) examine stability under learning of the targeted rational expectations equilibrium when policy-makers follow simple Taylor rules of various forms, including

$$i_t = k + \chi_\pi \pi_t + \chi_x x_t \quad (13)$$

$$i_t = k + \chi_\pi E_{t-1}^* \pi_t + \chi_x E_{t-1}^* x_t \text{ or} \quad (14)$$

$$i_t = k + \chi_\pi E_t^* \pi_{t+1} + \chi_x E_t^* x_{t+1}, \quad (15)$$

where $\chi_\pi, \chi_x \geq 0$ are policy parameters and the constant k reflects the steady state real interest rate and the target inflation rate. Bullard and Mitra (2002) show that under (13) and (14) the targeted rational expectations equilibrium is stable under learning if and only if

$$\lambda(\chi_\pi - 1) + (1 - \beta)\chi_x > 0, \quad (16)$$

a condition that holds if the Taylor principle $\chi_\pi > 1$ is satisfied. Under the forward-looking rule (15) this condition is still necessary, but there are additional stability conditions that require that χ_x and χ_π not be too large.¹⁸

Next, assume rational expectations for the moment and consider optimal policy obtained from minimizing a quadratic loss function

$$E_t \sum_{s=0}^{\infty} \beta^s (\pi_{t+s}^2 + \alpha x_{t+s}^2). \quad (17)$$

This type of optimal policy is often called “flexible inflation targeting”, see e.g. Svensson (1999) and Svensson (2003). α is the relative weight on the output target and pure inflation targeting would be the case $\alpha = 0$. The target for output is set at its efficient level. Without loss of generality for our purposes, the inflation target is set at zero. We treat the policy-maker’s preferences as exogenously given.¹⁹

The full intertemporal optimum under rational expectations, usually called the commitment solution, is obtained by maximizing (17) subject to (11)

¹⁸Additional considerations arise when there are heterogeneous expectations, see for example Honkapohja and Mitra (2005), Branch and McGough (2010), and Kurz (2011).

¹⁹It is well known, see Rotemberg and Woodford (1999), that the quadratic loss function (17) can be viewed as an approximation of the utility function of the representative consumer.

for all periods $t, t + 1, t + 2, \dots$. This solution leads to time inconsistency, and Woodford (1999a) and Woodford (1999b) have suggested that monetary policy ought be based on the *timeless perspective*. We refer to this as the “commitment solution” with the *commitment optimality condition*

$$\lambda\pi_t = -\alpha(x_t - x_{t-1}). \quad (18)$$

Assuming that agents are learning, Evans and Honkapohja (2003b) and Evans and Honkapohja (2006) consider optimal policy. For the commitment case Evans and Honkapohja (2006) show that an “expectations based” rule of the form

$$i_t = k + \delta_L x_{t-1} + \delta_\pi E_t^* \pi_{t+1} + \delta_x E_t^* x_{t+1} + \delta_g g_t + \delta_u u_t, \quad (19)$$

with coefficients chosen appropriated based on the structural parameters and the policy-maker loss function, can implement optimal policy, i.e. deliver an optimal rational expectations equilibrium that is stable under learning.

5.1.2 Policy and infinite-horizon learning

The infinite-horizon learning approach to monetary policy has been analyzed by Preston (2005) and Preston (2006). When Euler-equation learning is replaced by infinite-horizon learning the model becomes

$$x_t = E_t^* \sum_{T=t}^{\infty} \beta^{T-t} [(1 - \beta)x_{T+1} - \varphi(i_T - \pi_{T+1}) + g_T] \quad (20)$$

$$\pi_t = \lambda x_t + E_t^* \sum_{T=t}^{\infty} (\varpi\beta)^{T-t} [\lambda\varpi\beta x_{T+1} + (1 - \varpi)\beta\pi_{T+1} + u_T], \quad (21)$$

where ϖ is the Calvo stickiness parameter. On this approach the agents make fully optimal decisions, given their forecasts over the infinite future.

Both Euler-equation and infinite-horizon approaches are valid models of bounded rationality. The Euler-equation approach, which looks forward only one period, is clearly “boundedly optimal” in decisions and also boundedly rational in terms of forecasts because it does not explicitly incorporate long-term forecasts. On the other hand, in the infinite-horizon approach agents decisions depend on long-horizon forecasts, which will likely be modified by substantial amounts over time. As noted in Section 4.1, both types of learning can converge to the rational expectations equilibrium.

How does infinite-horizon learning affect the rational expectations equilibrium stability results for alternative interest-rate rules? Under infinite-horizon learning and the Taylor rule, condition (16) remains necessary but is

no longer sufficient for E-stability. Furthermore Preston (2006) argues that as $\beta \rightarrow 1$ E-stability cannot hold, under (14), and that it may not hold in calibrated models. The primary reason is that with long horizons, agents must forecast future interest rates as well as future inflation and the output gap. Indeed, if the private agents know the policy rule (14) and impose this relationship in their forecasts then (16) is again necessary and sufficient for stability. As Preston points out, these results can be interpreted as an argument for central bank communication.

A related point arises in connection with optimal policy. Evans and Honkapohja (2003b) and Evans and Honkapohja (2006) advocated “expectations based rules” with coefficients chosen to implement the first-order conditions for optimal policy. To implement this in the model (10)-(11) they recommended an interest rate obtained by solving (10), (11) and (18) simultaneously to eliminate π_t, x_t and obtain i_t in terms of expectations and fundamental shocks. This yields the rule (19), where $\delta_L = \frac{-\alpha}{\varphi(\alpha+\lambda^2)}$, $\delta_\pi = 1 + \frac{\lambda\beta}{\varphi(\alpha+\lambda^2)}$, $\delta_x = \delta_g = \varphi^{-1}$, $\delta_u = \frac{\lambda}{\varphi(\alpha+\lambda^2)}$. Evans and Honkapohja (2006) showed that this interest-rate rule guarantees determinacy and stability under learning.

If instead agents use long-horizon forecasts, then (19) can lead to instability under learning. In this case the expectations-based approach advocated by Evans and Honkapohja (2003b) and Evans and Honkapohja (2006) would have to be modified to use the behavioral equations (20)-(21). This is possible if observations of long-horizon expectations are available for the output gap, inflation and interest-rates, i.e.

$$\begin{aligned} ES_t^* z &= E_t^* \sum_{T=t}^{\infty} \beta^{T-t} z_{T+1}, \text{ for } z = x, \pi, i \\ E\tilde{S}_t^* z &= E_t^* \sum_{T=t}^{\infty} (\varpi\beta)^{T-t} z_{T+1} \text{ for } z = x, \pi. \end{aligned}$$

The long-horizon version of the Evans-Honkapohja expectations-based rule would then solve i_t from long-horizon IS and Phillips curves and the optimality condition (18). This yields a rule of the form

$$i_t = \psi_L x_{t-1} + \psi_\pi ES_t^* \pi + \psi_x ES_t^* x + \psi_i ES_t^* i + \tilde{\psi}_\pi E\tilde{S}_t^* \pi + \tilde{\psi}_x E\tilde{S}_t^* x + \psi_g g_t + \psi_u u_t,$$

where the coefficients are straightforward to compute. This rule would yield an optimal rational expectations equilibrium that is both determinate and stable with infinite-horizon learning.

As is clear from these results, the evolution of the economy under learning depends in part on the planning horizon of the agents. The interest-rate rule

used to implement optimal policy must therefore reflect the planning horizon as well as the rest of the economic structure.

5.2 Application: Price-Level Targeting and Optimal Policy

The research on learning and monetary policy has so far mostly considered the performance of interest-rate rules that implement inflation targeting under either discretion or commitment. There has been only limited research on the performance of price-level targeting from the viewpoint of learning stability and equilibrium determinacy.²⁰

The commitment optimality condition (18) can also be written in terms of the log of the price level p_t as $\lambda(p_t - p_{t-1}) = -\alpha(x_t - x_{t-1})$. This will be satisfied if

$$x_t = -\frac{\lambda}{\alpha}p_t + k, \quad (22)$$

for any constant k . Under price-level targeting the model with Euler-equation learning is given by the equations (10), (11) and (22), where we also specify that $\pi_t = p_t - p_{t-1}$.

We next compute the rational expectations equilibrium of interest using equations (11) and (22). Using the method of undetermined coefficients, the optimal rational expectations equilibrium can be expressed in the form

$$x_t = \bar{b}_p p_{t-1} + \bar{c}_x u_t + \bar{a}_x \quad (23)$$

$$p_t = \bar{b}_x p_{t-1} + \bar{c}_p u_t + \bar{a}_p. \quad (24)$$

Here $\bar{b}_x = (2\beta)^{-1}[\gamma - (\gamma^2 - 4\beta)^{1/2}]$ is the root inside the unit circle in the quadratic equation $\beta b_x^2 - \gamma b_x + 1 = 0$, where $\gamma = 1 + \beta + \lambda^2/\alpha$, and $\bar{b}_p = -\frac{\lambda}{\alpha}\bar{b}_x$, and the other unknown coefficients depend on the model parameters and the value of k .²¹

It is important to notice that the representation (23)-(24) of the optimal equilibrium does not indicate the form of the policy reaction function. Standard practice for rational expectations analysis is to calculate $E_t p_{t+1}$, $E_t x_{t+1}$ and to insert these and the rational expectations equilibrium p_t equation into

²⁰Evans and Honkapohja (2006), Gaspar, Smets, and Vestin (2007), and Preston (2008) consider aspects of price-level targeting under learning.

²¹The precise expressions for other coefficients are not needed in the analysis below.

the IS curve (10). This leads to the interest-rate reaction function

$$i_t = \eta_p p_{t-1} + \varphi^{-1} g_t + \eta_u u_t + \eta_0, \quad (25)$$

where $\eta_p = \bar{b}_x(1 - \bar{b}_x)(\frac{\lambda}{\alpha\varphi} - 1)$, $\eta_u = (1 - \bar{b}_x)(\frac{\lambda}{\alpha\varphi} - 1)\bar{c}_p$ and $\eta_0 = (\frac{\lambda}{\alpha\varphi} - 1)\bar{b}_x\bar{a}_p$. Equation (25) can be called a *fundamentals-based reaction function* as it indicates how to set optimally the policy instrument given the pre-determined and exogenous variables p_{t-1} , g_t and u_t .

We now examine the stability of both the fundamentals-based and an expectations-based interest-rate rule under learning.²² Substituting (25) into (10), the reduced form of the model can be written in the general form

$$y_t = A + ME_t^* y_{t+1} + QE_t^* y_t + Ny_{t-1} + Pv_t, \quad (26)$$

where $y_t = (x_t, p_t)'$ and the coefficients are given in Appendix 1. The analysis of learning stability for models of the general form (26) is developed in Appendix 1.

For the fundamentals-based interest-rate rule we have the following result:

Proposition 2 *Implementing price-level targeting using the fundamentals-based reaction function (25) does not guarantee local convergence of least-squares learning to the optimal rational expectations equilibrium. In particular, instability occurs if $\varphi > \frac{\lambda}{\alpha}$.*

It can be noted that, somewhat paradoxically, nearly strict inflation targeting (meaning that $\alpha \approx 0$) is very likely to lead to expectational instability.

An Expectations-Based Reaction Function The above computation deriving the fundamentals-based reaction function (25) relied heavily on the assumption that the economy is in the optimal rational expectations equilibrium. We now obtain the *expectations-based reaction function* under price-level targeting optimal policy. The policy rule is obtained by combining the optimality condition, the price-setting equation and the IS curve, for given private expectations.

Formally, we combine equations (10), (11), the optimality condition (22), and the definition of inflation in terms of current and price level $\pi_t = p_t - p_{t-1}$,

²²As noted above, we are here assuming short decision horizons based on the agents' Euler equations. Preston (2008) works out the analogous learning results for the model of price-level targeting with infinite-horizon decision rules.

treating private expectations as given.²³ Solving for i_t , the the expectations-based reaction function for price-level targeting is

$$i_t = \delta_L p_{t-1} + \delta_\pi E_t^* \pi_{t+1} + \delta_x E_t^* x_{t+1} + \delta_g g_t + \delta_u u_t, \quad (27)$$

where $\delta_L = \frac{\lambda}{\varphi(\alpha+\lambda^2)}$, $\delta_\pi = 1 + \frac{\lambda\beta}{\varphi(\alpha+\lambda^2)}$, $\delta_x = \varphi^{-1}$, $\delta_g = \varphi^{-1}$ and $\delta_u = \frac{\lambda}{\varphi(\alpha+\lambda^2)}$.

Looking at the rule (27) it can be seen that its coefficients stipulate a relatively large response to expected inflation ($\delta_\pi > 1$) and that effects coming from the expected output gap and the aggregate demand shock are fully neutralized ($\delta_x = \delta_g = \varphi^{-1}$). The next proposition shows that our interest-rate rule (27) performs well.

Proposition 3 *Implementing price-level targeting using the expectations-based reaction function (27), the optimal rational expectations equilibrium is stable under least-squares learning for all structural parameter values.*

Proposition 3 thus provides a remarkably strong result: Under the interest-rate setting rule (27), learning is stable and the economy is guided specifically to the optimal rational expectations equilibrium, and this result holds for all possible values of the structural parameters. Partial intuition for this result comes from the following (see also the reduced form (42) in Appendix 1). An increase in inflation expectations now leads to an increase in actual inflation that is smaller than the change in expectations. This dampened effect arises from the interest-rate reaction to changes in $E_t^* \pi_{t+1}$ and is a crucial element of the stability result.

5.3 Price-Level Targeting and the Zero Lower Bound

Discussions about the risk of deflation continue to be topical as the Western world is attempting to recover from the financial crisis. For example, see Bullard (2010). The experiences of 2008 and 2009 in several countries, as well as the experience of Japan since the 1990s, have created a situation in which the monetary policy response is constrained by the zero lower bound on nominal interest rates, a phenomenon sometimes called a ‘liquidity trap.’ Furthermore, in a liquidity trap there is the potential for the economy to get stuck in a deflationary situation with declining or persistently low levels of output.

²³Without loss of generality we set $k = 0$ in the formal analysis.

The possibility of the economy becoming trapped in a deflationary state, and the macroeconomic policies that might be able to avoid or extricate the economy from a liquidity trap, have been examined predominantly from the rational expectations perspective. Eggertsson and Woodford (2003) and Benhabib, Schmitt-Grohe, and Uribe (2001) show that besides the usual targeted steady state, there exists another steady state involving deflation. Eggertsson and Woodford (2003) argue that if the economy encounters a liquidity trap, monetary policy should commit to being expansionary for a considerable period of time, by keeping interest rates near zero even after the economy has emerged from deflation. Price-level targeting is, moreover, shown to be a nearly optimal policy response to the zero lower bound.²⁴

In our opinion, the rational expectations assumption is particularly questionable in an episode of deflation, which is far away from the inflation target and the normal state of the economy, and presents a new environment for economic agents. Our own view, reflected in Evans and Honkapohja (2005), Evans, Guse, and Honkapohja (2008) and Evans and Honkapohja (2010), is that the evolution of expectations plays a key role in the dynamics of the economy and that the tools from learning theory are needed for a realistic analysis of the zero lower bound issues.²⁵

These papers argued that when monetary policy is conducted using a Taylor rule with an inflation target, the usual targeted steady state is locally but not globally stable for least-squares and related learning processes. Moreover, deflation arises as a possible outcome when the economy is subject to large shocks. The workability of price-level targeting has, however, not been examined from the learning viewpoint, and as an application we take up this issue here.

Consider the following standard price-level targeting rule. It is like a Taylor rule, but instead of the inflation rate, the interest rate reacts to deviations from trend price-level target. These rules are sometimes called Wicksellian rules. The rule also takes into account the zero lower bound on interest rates in a piecewise-linear way.

²⁴Eggertsson and Woodford (2003) show that under rational expectations optimal policy can be implemented using a modified form of price-level targeting.

²⁵The first two papers formulate short-horizon learning using Euler equations, while the third paper postulates that agents use infinite-horizon decision rules. A related paper is Reifschneider and Williams (2000) that uses a fully backward-looking model.

The gross interest rate R_t follows the equation

$$R_t = 1 + \max[\bar{R} - 1 + \psi(p_t - \bar{p}_t), 0],$$

where trend inflation is modeled by $\bar{p}_t/\bar{p}_{t-1} = \pi^*$ and π^* is the inflation target. A steady state that satisfies $p_t = \bar{p}_t$ for all t , implies $R_t = \bar{R}$. If \bar{R} is set at the value $\bar{R} = \beta^{-1}\pi^*$, then the steady state also satisfies the Fisher equation, and the usual targeted steady state obtains under price-level targeting.

However, there is also a second steady state in which zero lower bound on interest rates is binding. If $p_t - \bar{p}_t \leq 0$ and $R_t = 1$, we can set

$$\pi_t = p_t/p_{t-1} = \beta$$

and obtain a second deflationary steady state due to a binding zero lower bound as the Fisher equation is also satisfied. It follows that, by itself, price-level targeting does not rule out a deflationary outcome.

Our objective is to consider the two steady states from the learning viewpoint. We employ a standard New Keynesian model with quadratic costs of price adjustment. For simplicity, we directly employ the model of Evans, Guse, and Honkapohja (2008), in which agents have short decision horizons with decisions given by the Euler equations for consumption and price-setting decisions. The system is

$$\beta \frac{\alpha\gamma}{\nu} (\pi_{t+1}^e - 1) \pi_{t+1}^e = -(c_t + g)^{(1+\varepsilon)/\alpha} + \frac{\alpha\gamma}{\nu} (\pi_t - 1) \pi_t \quad (28)$$

$$+ \alpha \left(1 - \frac{1}{\nu}\right) (c_t + g) c_t^{-\sigma_1}$$

$$c_t = c_{t+1}^e (\pi_{t+1}^e / \beta R_t)^{\sigma_1}, \quad (29)$$

$$R_t = \beta^{-1} + \psi(p_t - \bar{p}), \quad (30)$$

$$p_t = p_{t-1} \pi_t, \bar{p}_t = \bar{p}_{t-1} \pi^*, \quad (31)$$

where the first equation is the non-linear Phillips curve in implicit form. The second equation is the consumption Euler equation, while the third equation is the interest-rate rule. The last two equations define the actual and target price levels. One also adds the adjustment of expectations, which are

$$\pi_{t+1}^e = \pi_t^e + \phi(\pi_{t-1} - \pi_t^e)$$

$$c_{t+1}^e = c_t^e + \phi(c_{t-1} - c_t^e),$$

assuming constant-gain steady-state learning. ϕ is the gain parameter.

To simplify the analysis, assume that $\pi^* = 1$, so that $\bar{p}_t = \bar{p}$ for all t . Then the two steady states are the targeted steady state with a constant price level and the deflationary steady state $\pi_t = \beta$. We are interested in the possibility of local convergence of learning to each of the two steady states.

Consider the system near the targeted steady state $(c_t, \pi_t, p_t) = (c_H, 1, \bar{p})$. The analysis can be done using the nonlinear system (28)-(31) with point expectations.²⁶ Substituting the interest-rate rule into the consumption function yields an implicit system of three equations for the temporary equilibrium. Write this symbolically as

$$\mathcal{F}(y_t, y_{t+1}^e, y_{t-1}) = 0$$

using vector notation $y_t = (c_t, \pi_t, p_t)'$. Linearizing around the targeted steady state yields

$$y_t = (-D\mathcal{F}_y)^{-1}(DF_{y^e}y_{t+1}^e + D\mathcal{F}_{y_{-1}}y_{t-1}) \quad (32)$$

$$\equiv My_{t+1}^e + Ny_{t-1}, \quad (33)$$

where we have omitted constants and use the same notation for deviations from steady state. Appendix 2 works out the analytical details and shows the following sufficient condition for E-stability:

Proposition 4 *An active price-level targeting rule with $\psi > \beta^{-1}$ implies E-stability of the targeted steady state.*

For the deflationary steady state $(c_t, \pi_t, p_t) = (c_L, \beta, p_0\beta^t)$ the system consists of two key equations (28) and

$$c_t = c_{t+1}^e(\pi_{t+1}^e/\beta)^{\sigma_1},$$

since $R_t = 1$ near the deflationary steady state and p_t does not influence the dynamics of c_t and π_t . The system for $y_t' = (c_t, \pi_t)$ is purely forward-looking, so that the linearized system (32)-(33) has the form

$$y_t = (-D\mathcal{F}_y)^{-1}(D\mathcal{F}_{y^e})y_{t+1}^e = My_{t+1}^e,$$

and the perceived law of motion is just the intercept term. E-stability requires that the matrix $M - I$ is stable. Appendix 2 shows that we have the result (specified there more precisely):

²⁶Shocks have been dropped for brevity.

Proposition 5 *The deflationary steady state is usually not E-stable under price-level targeting.*

Moreover, since the interest rate is constant at $R = 1$ near the low steady state, the model is locally approximately the same as in Evans, Guse, and Honkapohja (2008). It follows that the deflationary steady state is a saddle-point in (π^e, c^e) space, implying that there exists a deflationary trap for the economy, in which there is worsening deflation and falling output over time.

6 Conclusions

This paper has addressed the question of whether there exist rational foundations for macroeconomics and finance models by focusing on the issue of how expectations are formed. This is a natural focus because of the centrality of “rational expectations” in modern macroeconomics, following the pioneering work of Muth (1961), Lucas (1972) and Sargent (1973). Expectations of future economic conditions are clearly a major determinant of current economic developments, and it can be argued that the role of expectations is the most fundamental way in which economics differs from the natural sciences.

How should we model how expectations are formed by households, firms, traders and policy-makers? The rational expectations hypothesis has a natural appeal to economists because it appears that optimizing agents would necessarily employ rational expectations – not to use rational expectations means that agents would be making systematic forecast errors, which must surely be to their disadvantage. For this reason rational expectations is necessarily the benchmark in macroeconomics.

But how can agents arrive at rational expectations? In the self-referential world of macroeconomics and finance, it is in fact rational to have rational expectations only if *other* agents have rational expectations. Rational expectations thus turns out to be an equilibrium concept. Once this is recognized, one is driven to consider the issue of stability: can agents plausibly coordinate on rational expectations? If this is possible, then it involves a process of learning in which agents with limited and imperfect knowledge come to hold rational expectations.

At this point we invoke the cognitive consistency principle that the agents in the economy should be assumed to be about as smart as are we, the economists studying the economy. This principle provides strong guidance, but

leads to several paths. For example, should we model agents like economic theorists, who think strategically about the actions of other agents within the context of a known structure? Or is it more plausible to model agents as applied econometricians trying to determine the appropriate specification of a forecasting model in a world of structural change? There appear to be diverse ways to model human, bounded, rationality and the appropriate choice presumably depends on the problem at hand. The range of plausible implementations can be regarded as a virtue, allowing for fruitful adaptation to a range of problems.

In this paper we have emphasized adaptive learning approaches, in which agents revise their forecasting model, or their choice of models, over time as new data become available. This approach has been fruitful in providing new approaches for explaining a range of phenomena, including the rise and fall of inflation, business cycle fluctuations, recurrent hyperinflations, and asset-price volatility and bubbles, and for studying the appropriate way to design monetary and fiscal policy.

To concretely illustrate the adaptive learning approach to expectation formation, we have focused on applications to monetary policy within New Keynesian models. The first application, in Section 5.1, to the choice of Taylor-type rules, stresses the potential importance, in models with private agent learning, of using monetary policies that do not simply react to the fundamental exogenous shocks, but which also explicitly react to private sector expectations. At the same time we found that it is important, when formulating the specific interest-rate rule, for policy-makers to understand the planning horizon used by private agents in making their decisions. In Section 5.2 we extend this analysis to interest-rate rules aiming to implement optimal policy using a form of price-level targeting. We again find that to ensure stability of the equilibrium under learning, it is desirable to condition policy on private-sector expectations.

Finally, in Section 5.3, we re-examine the zero lower bound constraint on interest rates. In previous work we have shown the possibility of a deflation trap region that can lead to unstable trajectories if the economy is subject to a large negative expectation shock. Extending this earlier work we consider whether a price-level targeting rule is sufficient to avoid this outcome. Our conclusion is that it is not, and that in some cases fiscal policy is needed to return the economy to the targeted steady state.

We think these applications illustrate the potential for adaptive learning approaches to provide a rational foundation for modeling expectations in

macroeconomics and finance.

Appendix 1

E-stability conditions in multivariate linear models.

We explain the main formal ideas for learning in the context of the general model (26). At each period t private agents have a perceived law of motion that they use to make forecasts about the future. The perceived law of motion takes the form

$$y_t = a_t + b_t y_{t-1} + c_t v_t, \quad (34)$$

where we are using the vector notation $y_t' = (x_t, p_t)$ and $v_t' = (g_t, u_t)$. For the reduced form (26) the optimal rational expectations equilibrium can be written as

$$y_t = \bar{a} + \bar{b} y_{t-1} + \bar{c} v_t,$$

where

$$\bar{a} = \begin{pmatrix} \bar{a}_x \\ \bar{a}_p \end{pmatrix}, \bar{b} = \begin{pmatrix} 0 & \bar{b}_x \\ 0 & \bar{b}_p \end{pmatrix} \text{ and } \bar{c} = \begin{pmatrix} 0 & \bar{c}_x \\ 0 & \bar{c}_p \end{pmatrix}.$$

Given the perceived law of motion and the current value of v_t , the forecast functions of the private agents are $E_t^* y_{t+1} = a_t + b_t E_t^* y_t + c_t E_t^* v_{t+1}$ or

$$E_t^* y_{t+1} = a_t + b_t(a_t + b_t y_{t-1} + c_t v_t) + c_t F v_t, \quad (35)$$

where (a_t, b_t, c_t) are the parameter values of the forecasts functions that agents have estimated on the basis of past data up to and including period $t - 1$. It is assumed that current exogenous variables, and lagged but not current endogenous variables, are in the information set when forecasts are made.

These forecasts are used in decisions for period t , which yields the actual law of motion for $y_t = (x_t, p_t)'$. The actual law of motion provides a new data point and agents are then assumed to re-estimate by least squares the parameters (a_t, b_t, c_t) with data through period t and use the updated forecast functions for period $t + 1$ decisions. Together with v_{t+1} these in turn yield the temporary equilibrium for period $t + 1$ and the learning dynamics continues with the same steps in subsequent periods. The rational expectations equilibrium $(\bar{a}, \bar{b}, \bar{c})$ is said to be *stable under least-squares learning* if the sequence (a_t, b_t, c_t) converges to $(\bar{a}, \bar{b}, \bar{c})$ over time.

Convergence conditions for least-squares learning can be obtained from E-stability conditions. For any model of the form (26), E-stability conditions are obtained as follows. Consider perceived law of motions of the form

$$y_t = a + b y_{t-1} + c v_t. \quad (36)$$

Computing expectations and inserting into (26) yields the actual law of motion

$$y_t = A + (Q + M(I + b))a + (Mb^2 + Qb + N)y_{t-1} + (Qc + M(bc + cF) + P)v_t. \quad (37)$$

This equation defines the mapping from perceived law of motion to actual law of motion parameters

$$T(a, b, c) = (A + (Q + M(I + b))a, Mb^2 + Qb + N, Qc + M(bc + cF) + P).$$

An rational expectations equilibrium $(\bar{a}, \bar{b}, \bar{c})$ is a fixed point of this map. Next, compute the derivative matrices

$$DT_a = Q + M(I + \bar{b}) \quad (38)$$

$$DT_b = \bar{b}' \otimes M + I \otimes M\bar{b} + I \otimes Q \quad (39)$$

$$DT_c = F' \otimes M + I \otimes M\bar{b} + I \otimes Q, \quad (40)$$

where \otimes denotes the Kronecker product and \bar{b} denotes the rational expectations equilibrium value of b . The necessary and sufficient condition for E-stability is that all eigenvalues of $DT_a - I$, $DT_b - I$ and $DT_c - I$ have negative real parts.²⁷

Proofs for Section 5.2.

Proof of Proposition 2: The coefficients of (26) are

$$A = \begin{pmatrix} -\varphi\eta_0 \\ -\lambda\varphi\eta_0 \end{pmatrix}, M = \begin{pmatrix} 1 & \varphi \\ \lambda & \beta + \lambda\varphi \end{pmatrix}, Q = \begin{pmatrix} 0 & -\varphi \\ 0 & -\beta - \lambda\varphi \end{pmatrix}, \quad (41)$$

$$N = \begin{pmatrix} 0 & -\varphi\eta_p \\ 0 & 1 - \lambda\varphi\eta_p \end{pmatrix} \text{ and } P = \begin{pmatrix} 0 & -\varphi\eta_u \\ 0 & 1 - \lambda\varphi\eta_u \end{pmatrix}.$$

We now apply the general E-stability conditions in Appendix 1 to the specific form of (26) with coefficients (41). A sufficient condition for instability is that one of the eigenvalues of $DT_a - I = Q + M(I + \bar{b}) - I$ has a positive real part. Hence a sufficient condition for instability is that $\det(DT_a - I) < 0$. It is easily computed that

$$Q + M(I + \bar{b}) - I = \begin{pmatrix} 0 & \bar{b}_p + \varphi\bar{b}_x \\ \lambda & \lambda\bar{b}_p + (\lambda\varphi + \beta)\bar{b}_x - 1 \end{pmatrix}$$

²⁷We are excluding the exceptional cases where one or more eigenvalue has zero real part.

and $\det(DT_a - I) = -\lambda(\bar{b}_p + \varphi\bar{b}_x)$. Using $\bar{b}_p = -\frac{\lambda}{\alpha}\bar{b}_x$ we get $\det(DT_a - I) = -\lambda(\varphi - \frac{\lambda}{\alpha})\bar{b}_x$. Since $0 < \bar{b}_x < 1$, this is negative when $\varphi > \frac{\lambda}{\alpha}$. Q.E.D.

Proof of Proposition 3: The reduced form of the economy under (27) is given in Appendix 1.

$$\begin{pmatrix} x_t \\ \pi_t \\ p_t \end{pmatrix} = \begin{pmatrix} 0 & -\frac{\lambda\beta}{\alpha+\lambda^2} & 0 \\ 0 & \frac{\alpha\beta}{\alpha+\lambda^2} & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} E_t^* x_{t+1} \\ E_t^* \pi_{t+1} \\ E_t^* p_{t+1} \end{pmatrix} + \begin{pmatrix} \frac{\alpha}{\alpha+\lambda^2} & 0 & 0 \\ \frac{\alpha\lambda}{\alpha+\lambda^2} & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} x_{t-1} \\ \pi_{t-1} \\ p_{t-1} \end{pmatrix} + \begin{pmatrix} 0 & -\frac{\lambda}{\alpha+\lambda^2} \\ 0 & \frac{\alpha}{\alpha+\lambda^2} \\ 0 & 0 \end{pmatrix} u_t. \quad (42)$$

This is thus in standard form. Using Mathematica,²⁸ it can be computed that the eigenvalues of $DT_a - I$ are two repeated roots of -1 and one root equal to $-1 + \frac{\alpha\beta}{\alpha+\lambda^2}$. Those of $DT_b - I$ are eight repeated roots of -1 and one root equal to $-1 + \frac{\alpha\beta}{\alpha+\lambda^2}$. Those of $DT_c - I$ are four roots of -1 , $-1 + \frac{\alpha\beta\mu}{\alpha+\lambda^2}$ and $-1 + \frac{\alpha\beta\rho}{\alpha+\lambda^2}$. Thus all eigenvalues are negative, so that learning stability obtains. Q.E.D.

Appendix 2

E-stability in the model of Section 5.3. The E-stability conditions in Appendix 1 require that all eigenvalues of the matrices $M(I + \bar{b})$ and $\bar{b}' \otimes M + I \otimes M\bar{b}$ have real parts less than one. The general case leads to complex expressions, but in the limit $\gamma \rightarrow 0$ one obtains using Mathematica

$$M = \begin{pmatrix} 0 & 0 & 0 \\ \frac{\sigma_1}{c_H \bar{p} \beta \psi} & \frac{1}{\bar{p} \beta \psi} & 0 \\ \frac{\sigma_1}{c_H \beta \psi} & \frac{1}{\beta \psi} & 0 \end{pmatrix}, \quad N = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & \frac{-1}{\bar{p}} \\ 0 & 0 & 0 \end{pmatrix}.$$

From now on, we normalize $\bar{p} = 1$.

Assuming that the solution takes the form $y_t = a + by_{t-1}$ and solving the equation $Mb^2 - b + N = 0$, one obtains

$$a = 0, \quad b = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & b_{23} \\ 0 & 0 & b_{33} \end{pmatrix},$$

²⁸The Mathematica routines used in the paper are available on request.

where $(\beta\psi)^{-1}b_{23}^2 - b_{23} - 1 = 0$ and $b_{33} = (\beta\psi)^{-1}b_{23}^2$. The quadratic for b_{23} has two roots, from which we pick the negative root

$$b_{23} = \frac{1}{2} \left(\beta\psi - \sqrt{\beta\psi} \sqrt{4 + \beta\psi} \right).$$

The E-stability conditions are that the matrices $DT_a - I = M(I + b) - I$ and $DT_b - I = \bar{b}' \otimes M + I \otimes M\bar{b} - I$ are stable, i.e. all of their eigenvalues have real parts less than one. Two of the eigenvalues of $DT_a - I$ are -1 and the third one is

$$\frac{2 - \beta\psi - \sqrt{\beta\psi} \sqrt{4 + \beta\psi}}{2\beta\psi}.$$

The matrix $DT_b - I$ have six roots of -1 , two roots equal to

$$-1 + \frac{\beta\psi - \sqrt{\beta\psi} \sqrt{4 + \beta\psi}}{2\beta\psi}$$

and the remaining root equal to

$$\frac{1 - \beta\psi - \sqrt{\beta\psi} \sqrt{4 + \beta\psi}}{\beta\psi}.$$

Clearly, for stability it is sufficient that $\beta\psi > 1$ or $\psi > \beta^{-1}$.

For the deflationary steady state, the determinant of the E-stability system is equal to

$$\frac{-c_L^{1+\sigma_1} (c_L + g)^{\frac{1+\varepsilon}{\alpha}} (1 + \varepsilon) \nu - g^2 \alpha^2 (\nu - 1) \sigma_1 - c_L \alpha^2 (\nu - 1) [g(2\sigma_1 - 1) + c_L(\sigma_1 - 1)]}{c_L^{\sigma_1} ((c_L + g) \alpha^2 \beta (2\beta - 1) \gamma \sigma_1)}.$$

It is certainly negative for $\sigma_1 \geq 1$ and even when a weaker (but cumbersome) condition is satisfied.

References

- ADAM, K. (2007): “Experimental Evidence on the Persistence of Output and Inflation,” *Economic Journal*, 117, 603–636.
- ADAM, K., A. MARCET, AND J. P. NICOLINI (2008): “Stock Market Volatility and Learning,” Working paper.
- BARRO, R. J. (1989): “The Ricardian Approach to Budget Deficits,” *Journal of Economic Perspectives*, 3, 37–54.
- BASDEVANT, O. (2005): “Learning Processes and Rational Expectations: an Analysis Using a Small Macro-econometric Model for New Zealand,” *Economic Modelling*, 22, 1074–1089.
- BENHABIB, J., S. SCHMITT-GROHE, AND M. URIBE (2001): “The Perils of Taylor Rules,” *Journal of Economic Theory*, 96, 40–69.
- BERARDI, M. (2007): “Heterogeneity and Misspecification in Learning,” *Journal of Economic Dynamics and Control*, 31, 3203–3227.
- (2008): “Monetary Policy with Heterogeneous and Misspecified Expectations,” *Journal of Money, Credit, and Banking*, 41, 79–100.
- BERNANKE, B., AND M. WOODFORD (eds.) (2005): *The Inflation-Targeting Debate*. University of Chicago Press, Chicago, USA.
- BERNHEIM, B. D. (1987): “Ricardian Equivalence: An Evaluation of Theory and Evidence,” *NBER Macroeconomics Annual*, 2, 263–304.
- BRANCH, W. A. (2004): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *Economic Journal*, 114, 592–621.
- BRANCH, W. A., AND G. W. EVANS (2006a): “Intrinsic Heterogeneity in Expectation Formation,” *Journal of Economic Theory*, 127, 264–295.
- (2006b): “A Simple Recursive Forecasting Model,” *Economic Letters*, 91, 158–166.
- (2007): “Model Uncertainty and Endogenous Volatility,” *Review of Economic Dynamics*, 10, 207–237.

- (2010): “Asset Return Dynamics and Learning,” *Review of Financial Studies*, 23, 1651–1680.
- (2011): “Learning about Risk and Return: A Simple Model of Bubbles and Crashes,” *American Economic Journal: Macroeconomics*, 3, 159–191.
- BRANCH, W. A., G. W. EVANS, AND B. MCGOUGH (2010): “Finite Horizon Learning,” Working paper.
- BRANCH, W. A., AND B. MCGOUGH (2010): “Dynamic Predictor Selection in a New Keynesian Model with Heterogeneous Expectations,” *Journal of Economic Dynamics and Control*, 34, 1492–1508.
- (2011): “Heterogeneous Expectations, Shock Amplification and Complex Dynamics in Competitive Business Cycle Models,” *Economic Theory*, 47, 395–422.
- BRAY, M., AND N. SAVIN (1986): “Rational Expectations Equilibria, Learning, and Model Specification,” *Econometrica*, 54, 1129–1160.
- BRAZIER, A., R. HARRISON, M. KING, AND T. YATES (2008): “The Danger of Inflating Expectations of Macroeconomic Stability: Heuristic Switching in an Overlapping Generations Monetary Model,” *International Journal of Central Banking*, 4, 219–254.
- BROCK, W., C. HOMMES, AND F. WAGENER (2009): “More Hedging Instruments May Destabilize Markets,” *Journal of Economic Dynamics and Control*, 33, 1912–1928.
- BROCK, W. A., AND C. H. HOMMES (1997): “A Rational Route to Randomness,” *Econometrica*, 65, 1059–1095.
- (1998): “Heterogeneous Beliefs and Routes to Chaos in a Simple Asset pricing Model,” *Journal of Economic Dynamics and Control*, 22, 1235–1274.
- BULLARD, J. (2006): “The Learnability Criterion and Monetary Policy,” *Federal Reserve Bank of St. Louis Review*, 88, 203–217.
- (2010): “Seven Faces of The Peril,” *Federal Reserve Bank of St. Louis Review*, 92, 339–352.

- BULLARD, J., AND S. EUSEPI (2005): “Did the Great Inflation Occur Despite Policymaker Commitment to a Taylor Rule?,” *Review of Economic Dynamics*, 8, 324–359.
- BULLARD, J., AND K. MITRA (2002): “Learning About Monetary Policy Rules,” *Journal of Monetary Economics*, 49, 1105–1129.
- CARBONI, G., AND M. ELLISON (2008): “The Great Inflation and the Greenbook,” mimeo.
- CHAKRABORTY, A., AND G. W. EVANS (2008): “Can Perpetual Learning Explain the Forward Premium Puzzle?,” *Journal of Monetary Economics*, 55, 477–490.
- CHO, I.-K., AND K. KASA (2010): “Learning and Model Validation: An Example,” mimeo.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): “Escaping Nash Inflation,” *Review of Economic Studies*, 69, 1–40.
- CLARIDA, R., J. GALI, AND M. GERTLER (1999): “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37, 1661–1707.
- COBHAM, D., Ø. EITRHEIM, S. GERLACH, AND J. F. QVIGSTAD (eds.) (2010): *Twenty Years of Inflation Targeting: Lessons Learned and Future Prospects*. Cambridge University Press, Cambridge.
- COGLEY, T., AND T. J. SARGENT (2005): “The Conquest of US Inflation: Learning and Robustness to Model Uncertainty,” *Review of Economic Dynamics*, 8, 528–563.
- (2008): “Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making,” *International Economic Review*, 49, 185–221.
- DE GRAUWE, P. (2011): “Animal Spirits and Monetary Policy,” *Economic Theory*, forthcoming.
- DE GRAUWE, P., AND M. GRIMALDI (2006): *The Exchange Rate in a Behavioral Finance Framework*. Princeton University Press, Princeton.

- EGGERTSSON, G. B., AND M. WOODFORD (2003): “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity*, (1), 139–233.
- ELLISON, M., AND T. YATES (2007): “Escaping Volatile Inflation,” *Journal of Money, Credit and Banking*, 39, 981–993.
- EUSEPI, S., AND B. PRESTON (2007): “Central Bank Communication and Expectations Stabilization,” Working paper nr.13259, NBER.
- (2011): “Expectations, Learning and Business Cycle Fluctuations,” *American Economic Review*, forthcoming.
- EVANS, G. W., AND R. GUESNERIE (1993): “Rationalizability, Strong Rationality, and Expectational Stability,” *Games and Economic Behaviour*, 5, 632–646.
- (2003): “Coordination on Saddle Path Solutions: the Eductive Viewpoint – Linear Univariate Models,” *Macroeconomic Dynamics*, 7, 42–62.
- EVANS, G. W., R. GUESNERIE, AND B. MCGOUGH (2010): “Eductive Stability in Real Business Cycle Models,” mimeo.
- EVANS, G. W., E. GUSE, AND S. HONKAPOHJA (2008): “Liquidity Traps, Learning and Stagnation,” *European Economic Review*, 52, 1438–1463.
- EVANS, G. W., AND S. HONKAPOHJA (1996): “Least Squares Learning with Heterogeneous Expectations,” *Economics Letters*, 53, 197–201.
- (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.
- (2003a): “Adaptive Learning and Monetary Policy Design,” *Journal of Money, Credit and Banking*, 35, 1045–1072.
- (2003b): “Expectations and the Stability Problem for Optimal Monetary Policies,” *Review of Economic Studies*, 70, 807–824.
- (2005): “Policy Interaction, Expectations and the Liquidity Trap,” *Review of Economic Dynamics*, 8, 303–323.

- (2006): “Monetary Policy, Expectations and Commitment,” *Scandinavian Journal of Economics*, 108, 15–38.
- (2009a): “Expectations, Learning and Monetary Policy: An Overview of Recent Research,” in Schmidt-Hebbel and Walsh (2009), chap. 2, pp. 27–76.
- (2009b): “Learning and Macroeconomics,” *Annual Review of Economics*, 1, 421–451.
- (2010): “Expectations, Deflation Traps and Macroeconomic Policy,” in Cobham, Eitrheim, Gerlach, and Qvigstad (2010), chap. 12, pp. 232–260.
- EVANS, G. W., S. HONKAPOHJA, AND R. MARIMON (2001): “Convergence in Monetary Inflation Models with Heterogeneous Learning Rules,” *Macroeconomic Dynamics*, 5, 1–31.
- (2007): “Stable Sunspot Equilibria in a Cash-in-Advance Economy,” *The B.E. Journal of Macroeconomics (Advances)*, 7, Iss. 1, Article 3.
- EVANS, G. W., S. HONKAPOHJA, AND K. MITRA (2009): “Anticipated Fiscal Policy and Learning,” *Journal of Monetary Economics*, 56, 930–953.
- (2010): “Does Ricardian Equivalence Hold When Expectations are not Rational?” mimeo.
- EVANS, G. W., S. HONKAPOHJA, AND P. ROMER (1998): “Growth Cycles,” *American Economic Review*, 88, 495–515.
- EVANS, G. W., AND B. MCGOUGH (2010): “Learning to Optimize,” mimeo.
- EVANS, G. W., AND G. RAMEY (1992): “Expectation Calculation and Macroeconomic Dynamics,” *American Economic Review*, 82, 207–224.
- (1998): “Calculation, Adaptation and Rational Expectations,” *Macroeconomic Dynamics*, 2, 156–182.
- FEDERAL RESERVE BANK, K. C. (1999): *New Challenges for Monetary Policy*. Federal Reserve Bank of Kansas City, Kansas City.

- FOURGEAUD, C., C. GOURIEROUX, AND J. PRADEL (1986): “Learning Procedures and Convergence to Rationality,” *Econometrica*, 54, 845–868.
- FRYDMAN, R., AND M. D. GOLDBERG (2010a): “The Imperfect Knowledge Imperative in Modern Macroeconomics and Finance Theory,” mimeo.
- (2010b): “Opening Models of Asset Prices and Risk to Non-Routine Change,” mimeo.
- FUSTER, A., D. LAIBSON, AND B. MENDEL (2010): “Natural Expectations and Macroeconomic Fluctuations,” *Journal of Economic Perspectives*, 24, 67–84.
- GASPAR, V., F. SMETS, AND D. VESTIN (2007): “Is Time ripe for Price Level Path Stability?,” Working paper no 818, European Central Bank.
- GIANNITSAROU, C. (2003): “Heterogenous Learning,” *Review of Economic Dynamics*, 6, 885–906.
- (2006): “Supply-Side Reforms and Learning Dynamics,” *Journal of Monetary Economics*, 53, 291–309.
- GUESNERIE, R. (1992): “An Exploration of the Eductive Justifications of the Rational-Expectations Hypothesis,” *American Economic Review*, 82, 1254–1278.
- (2002): “Anchoring Economic Predictions in Common Knowledge,” *Econometrica*, 70, 439–480.
- GUESNERIE, R., AND J.-C. ROCHET (1993): “(De)stabilizing Speculation on Futures Markets: An Alternative Viewpoint,” *European Economic Review*, 37, 1043–1063.
- GUSE, E. (2006): “Stability Properties for Learning with Heterogeneous Expectations and Multiple Equilibria,” *Journal of Economic Dynamics and Control*, 29, 1623–1642.
- HENS, T., AND K. SHENK-HOPPÉ (eds.) (2009): *Handbook on Financial Markets: Dynamics and Evolution*. North Holland.

- HOMMES, C. H. (2011): “The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab,” *Journal of Economic Dynamics and Control*, 35, 1–24.
- HOMMES, C. H., AND G. SORGER (1997): “Consistent Expectations Equilibria,” *Macroeconomic Dynamics*, 2, 287–321.
- HONKAPOHJA, S., AND K. MITRA (2005): “Performance of Monetary Policy with Internal Central Bank Forecasting,” *Journal of Economic Dynamics and Control*, 29, 627–658.
- (2006): “Learning Stability in Economies with Heterogeneous Agents,” *Review of Economic Dynamics*, 9, 284–309.
- HONKAPOHJA, S., K. MITRA, AND G. W. EVANS (2002): “Notes on Agents’ Behavioral Rules Under Adaptive Learning and Recent Studies of Monetary Policy,” manuscript, www.valt.helsinki.fi/RAKA/seppo.html, and <http://darkwing.uoregon.edu/~gevans>.
- HUANG, K., Z. LIU, AND T. ZHA (2009): “Learning, Adaptive Expectations and Technology Shocks,” *Economic Journal*, 119, 377–405.
- JACOBS, D., E. KALAI, AND M. KAMIEN (eds.) (1998): *Frontiers of Research in Economic Theory*. Cambridge University Press, Cambridge.
- JUSELIUS, K. (2010): “Imperfect Knowledge, Asset Price Swings and Structural Slumps: A Cointegrated VAR Analysis of their Interdependence,” mimeo.
- KASA, K. (2004): “Learning, Large Deviations, and Recurrent Currency Crises,” *International Economic Review*, 45, 141–173.
- KIM, Y. S. (2009): “Exchange Rates and Fundamentals under Adaptive Learning,” *Journal of Economic Dynamics and Control*, 33, 843–863.
- KREPS, D. M. (1998): “Anticipated Utility and Dynamic Choice,” in Jacobs, Kalai, and Kamien (1998), pp. 242–274.
- KURZ, M. (ed.) (1997): *Endogenous Economic Fluctuations. Studies in the Theory of Rational Beliefs*. Springer-Verlag, Berlin.

- (2009): “Rational Diverse Beliefs and Economic Volatility,” in Hens and Shenk-Hoppé (2009), chap. 8, pp. 439–506.
- (2011): “A New Keynesian Model with Diverse Beliefs,” mimeo.
- LANSING, K. (2010): “Rational and Near-rational Bubbles without Drift,” *Economic Journal*, 120, 1149–1174.
- LEBARON, B. (2010): “Heterogeneous Gain Learning and Long Swings in Asset Prices,” mimeo.
- LUCAS, JR., R. E. (1972): “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, 4, 103–124.
- (1973): “Some International Evidence on Output-Inflation Trade-offs,” *American Economic Review*, 63, 326–334.
- MARCET, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93, 1476–1498.
- MARCET, A., AND T. J. SARGENT (1989a): “Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information,” *Journal of Political Economy*, 97, 1306–1322.
- (1989b): “Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models,” *Journal of Economic Theory*, 48, 337–368.
- MARIMON, R., AND S. SUNDER (1993): “Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence,” *Econometrica*, 61, 1073–1107.
- (1995): “Does a Constant Money Growth Rule Help Stabilize Inflation?: Experimental Evidence,” *Carnegie-Rochester Conference Series on Public Policy*, 43, 111–156.
- MARK, N. (2007): “Changing Monetary Rules, Learning, and Real Exchange Rate Dynamics,” manuscript.
- MARKIEWICZ, A. (2010): “Monetary Policy, Model Uncertainty, and Exchange Rate Volatility,” manuscript.

- MILANI, F. (2007): “Expectations, Learning and Macroeconomic Persistence,” *Journal Of Monetary Economics*, 54, 2065–2082.
- (2011): “Expectation Shocks and Learning as Drivers of the Business Cycle,” *Economic Journal*, 121, 379–401.
- MUTH, J. F. (1961): “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29, 315–335.
- ORPHANIDES, A., AND J. C. WILLIAMS (2005a): “The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning and Expectations,” *Journal of Economic Dynamics and Control*, 29, 1927–1950.
- (2005b): “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in Bernanke and Woodford (2005), chap. 5, pp. 201–234.
- (2005c): “Inflation Scares and Forecast-Based Monetary Policy,” *Review of Economic Dynamics*, 8, 498–527.
- (2007): “Robust Monetary Policy with Imperfect Knowledge,” *Journal of Monetary Economics*, 54, 1406–1435.
- PFAJFAR, D. (2007): “Formation of Rationally Heterogeneous Expectations,” mimeo.
- PFAJFAR, D., AND E. SANTORO (2007): “Heterogeneity and Learning in Inflation Expectation Formation: An Empirical Assessment,” mimeo.
- PRESTON, B. (2005): “Learning about Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1, 81–126.
- (2006): “Adaptive Learning, Forecast-based Instrument Rules and Monetary Policy,” *Journal of Monetary Economics*, 53, 507–535.
- (2008): “Adaptive Learning, Forecast-based Instrument Rules and Monetary Policy,” *Journal of Economic Dynamics and Control*, 32, 3661–3681.
- PRIMICERI, G. E. (2006): “Why Inflation Rose and Fell: Policy-Makers’ Beliefs and U. S. Postwar Stabilization Policy,” *Quarterly Journal of Economics*, 121, 867–901.

- REIFSCHEIDER, D., AND J. C. WILLIAMS (2000): “Three Lessons for Monetary Policy in a Low-Inflation Era,” *Journal of Money, Credit and Banking*, 32, 936–966.
- RICCIUTI, R. (2003): “Assessing Ricardian Equivalence,” *Journal of Economic Surveys*, 17, 55–78.
- ROTEMBERG, J. J., AND M. WOODFORD (1999): “Interest Rate rules in an Estimated Sticky-Price Model,” in Taylor (1999), pp. 57–119.
- SARGENT, T. J. (1973): “Rational Expectations, the Real Rate of Interest and the Natural Rate of Unemployment,” *Brookings Papers on Economic Activity*, 2, 429–472.
- (1991): “Equilibrium with Signal Extraction from Endogenous Variables,” *Journal of Economic Dynamics and Control*, 15, 245–273.
- (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.
- (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton NJ.
- SARGENT, T. J., AND N. WALLACE (1973a): “Rational Expectations and the Dynamics of Hyperinflation,” *International Economic Review*, 14, 429–472.
- (1973b): “The Stability of Models of Money and Growth with Perfect Foresight,” *Econometrica*, 41, 1043–1048.
- SARGENT, T. J., N. WILLIAMS, AND T. ZHA (2006): “Shocks and Government Beliefs: The Rise and Fall of American Inflation,” *American Economic Review*, 96, 1193–1224.
- SCHMIDT-HEBBEL, K., AND C. E. WALSH (2009): *Monetary Policy under Uncertainty and Learning*. Central Bank of Chile, Santiago.
- SEATER, J. J. (1993): “Ricardian Equivalence,” *Journal of Economic Literature*, 31, 142–190.
- SLOBODYAN, S., AND R. WOUTERS (2007): “Learning in an Estimated DSGE Model,” mimeo.

- (2008): “Estimating a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models,” mimeo.
- SVENSSON, L. E. (1999): “Inflation Targeting as a Monetary Policy Rule,” *Journal of Monetary Economics*, 43, 607–654.
- (2003): “What is Wrong with Taylor Rules? Using Judgement in Monetary Policy through Targeting Rules,” *Journal of Economic Literature*, 41, 426–477.
- TAYLOR, J. (ed.) (1999): *Monetary Policy Rules*. University of Chicago Press, Chicago.
- TIMMERMANN, A. G. (1993): “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices,” *Quarterly Journal of Economics*, 108, 1135–1145.
- (1996): “Excessive Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning,” *Review of Economic Studies*, 63, 523–557.
- VAN NIEUWERBURGH, S., AND L. VELDKAMP (2006): “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics*, 53, 753–772.
- WILLIAMS, N. (2004): “Adaptive Learning and Business Cycles,” working paper.
- WOODFORD, M. (1999a): “Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?,” in Federal Reserve Bank (1999), pp. 277–316.
- (1999b): “Optimal Monetary Policy Inertia,” *The Manchester School, Supplement*, 67, 1–35.