Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

A Short History in Defence of Adaptive Learning

CAMA Working Paper 52/2023
October 2023

Stefano Eusepi
University of Texas

Bruce Preston
University of Melbourne
Centre for Applied Macroeconomic Analysis, ANU

Abstract

This paper tells the story of adaptive learning. With origins in the rational expectations revolution, models of adaptive learning are as much about moving people's expectations closer to rational expectations as further away. But relaxing the rational expectations requirement that subjective expectations are model consistent permits addressing a broader class of economic questions as well as resolving certain empirical puzzles. Importantly, models of optimal intertemporal decision making with adaptive learning are supported by a growing body of empirical evidence using surveys of professional, household, firm and policymaker forecasts.

Keywords

adaptive expectations, rational expectations, adaptive learning, inflation expectations

JEL Classification

E32, D83, D84

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

The Centre for Applied Macroeconomic Analysis in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

A Short History in Defence of Adaptive Learning*

Stefano Eusepi† Bruce Preston‡

Abstract

This paper tells the story of adaptive learning. With origins in the rational expectations revolution, models of adaptive learning are as much about moving people's expectations closer to rational expectations as further away. But relaxing the rational expectations requirement that subjective expectations are model consistent permits addressing a broader class of economic questions as well as resolving certain empirical puzzles. Importantly, models of optimal intertemporal decision making with adaptive learning are supported by a growing body of empirical evidence using surveys of professional, household, firm and policymaker forecasts.

Keywords: Adaptive expectations, rational expectations, adaptive learning, inflation expectations
JEL Codes: E32, D83, D84

*October 18, 2023. The authors thank David Byrne, Chris Gibbs, Greg Kaplan and Michael Woodford for comments. The authors acknowledge research support from the Australian Research Council in the grant DP210103427. The usual caveat applies.

†University of Texas, Austin. E-mail: stefano.eusepi@austin.utexas.edu.

‡The University of Melbourne and CAMA. E-mail: bruce.preston@unimelb.edu.au.

1 Introduction

Contemporary discussions of adaptive learning are framed by the "totally victorious" rational expectations revolution.[1] In modern macroeconomic models rational expectations equilibrium is premised on two stipulations: optimization and mutual consistency of beliefs, the idea that subjective and objective beliefs coincide, so that expectations are consistent with the relevant economic theory. Adaptive learning is then cast as a retreat from this second stipulation, a reversion to the adaptive expectations of old and therefore subject to familiar criticism.

But the story of adaptive learning is as much about building theories of expectations that move people's expectations closer to rational expectations as they are about moving further away. Having antecedents in the adaptive expectations literature, the adaptive learning literature crystallized as part of the rational expectations revolution. While presenting an elegant solution to the problem of expectations modeling in intertemporal economic models, rational expectations also raised questions of its own. To be plausible surely we must explain how people come to hold such expectations? And because rational expectations theory was rife with equilibrium indeterminacy, surely we must explain how people coordinate on one particular equilibrium. By providing a behavioral theory of expectations formation, adaptive learning models offered possible justification of rational expectations equilibrium and also an equilibrium selection mechanism. But it also held the promise to deliver much more in the form of an economic theory of transition. How do people come to understand the consequences of new policies and significant structural changes to the economy?

To varying degrees this literature made progress on these and other questions. Perhaps to the detriment of its own success, the answers were not decisively in favor of rational expectations equilibrium being the natural outcome of some adaptive learning process. The problem was not so much that people could not learn a rational expectations equilibrium—though this was certainly true of some models—but that learning did not always provide a clear equilibrium selection criterion. Models predicted that learning could converge to sunspot equilibria and even cycles between such equilibria. And while modern models for monetary policy analysis are less susceptible to these concerns—as policy can be chosen appropriately to ensure convergence to desirable equilibria—there was a sense that anything was possible.

Over time, perhaps because of the absence of a compelling and unifying theory of how a specific rational expectations equilibrium would arise, the urgency to understand the knowledge foundations of rational expectations diminished, and in some intellectual quarters com-

[1] See the interview with Robert Lucas in Bowmaker (2013).

pletely vanished. And as the learning literature branched out to explore quantitative and empirical problems, the criticisms of adaptive expectations were applied anew: that beliefs were not 'optimal'; there are too many degrees of freedom making the theory untestable; that rational people would never make systematic forecast errors; that such beliefs fail to predict how expectations change given changes in policy. But these criticisms, aside from being misplaced, misunderstand the intent of the adaptive learning literature. The approach is as much concerned with placing a lower bound on rationality as it is about placing an upper bound in acknowledgment of the complexity of economic decision making.

At the heart of new theories of adaptive learning is the question of what it means to be rational. Benjamin Friedman (1979) cleverly addresses this issue using words from John Rawls in his *Theory of Justice*

> One might reply that the rationality of a person's choice does not depend upon how much he knows, but only upon how well he reasons from whatever information he has, however incomplete. Our decision is perfectly rational provided that we face up to our circumstances and do the best we can.

Adaptive agents make optimal decisions and process their information efficiently, just like econometricians would do. But their subjective beliefs need not coincide with objective beliefs, as their forecasting models are based on incomplete knowledge about a complex economic environment that is constantly changing.

The critical task of adaptive learning research then is to discipline deviations from the full rationality benchmark placing a lower bound on rationality. One approach chooses learning mechanisms that, in equilibrium, respect pre-defined bounds to rationality. The idea being that 'large' discrepancies between subjective and objective expectations would induce individual agents to deviate from the proposed equilibrium. A second, possibly complementary approach, uses observed expectations from surveys to identify learning mechanisms and to measure the discrepancy between subjective and objective beliefs. Growing data availability over the past thirty years has given rise to a large literature using survey data from professional forecasters, households, firms and policymakers to test rationality of expectations and select alternative expectations mechanisms. The evidence challenges the full information rational expectations paradigm. And despite an active debate about alternative theories of information frictions, a new dominant paradigm has yet to emerge. This short history 'makes the case' for models of adaptive learning both on theoretical and empirical grounds.

This essay is organized as follows. Section 2 develops a short history of adaptive learning, starting from contributions in the early 1950s, covering the rational expectations revolution and subsequent adaptive learning literature. Section 3 discusses recent methodological advances which permit modeling adaptive learning with optimal decisionmaking. Section 4 provides three examples which demonstrates optimal decision making under adaptive learning provides a compelling framework to answer a range of economic questions, particularly those relating to policy uncertainty and regime change. Section 5 briefly discusses empirical work using survey data on expectations that adduces clear evidence in support of models of adaptive learning. Section 6 concludes.

2 A Short History of Adaptive Learning

This is the story of adaptive learning. As an organizing framework, we focus on inflation dynamics and the so-called Phillips curve, which as a theory of inflation determination was central to the development of many of the ideas we cover.

2.1 Adaptive Expectations: 1950s

Our journey commences with Arrow and Nerlove's (1958) reinterpretation of Hick's (1939) definition of the "elasticity of expectations"

$$\hat{E}_t\pi_{t+1} = \hat{E}_{t-1}\pi_t + g\left(\pi_t - \hat{E}_{t-1}\pi_t\right) \tag{1}$$

where the operator $\hat{E}_t$ denotes expectations of future inflation π_{t+1} formed in period t; and the parameter g, the elasticity, determines the degree to which expectations are adjusted in response to the discrepancy between the actual and expected state of affairs, formally the forecast error. When $g = 1$ the current inflation rate is taken as the expectation of future inflation, sometimes called static expectations; when $g = 0$ current inflation has no effect on expected inflation. Intermediate values permit more general relationships between forecast errors and the adjustment of expectations.

This theory of belief formation is called adaptive expectations. The theory was central to concerns about the dynamic stability of economies. Kaldor (1934) used static expectations to establish the "cobweb theorem". More general applications followed, including Cagan's (1954) study of hyperinflations; Friedman's (1957) study of permanent income theory; and Nerlove's (1958) study of agricultural cycles.

For the time being, we will take the elasticity as given, consistent with this research. We later show that modern analysis of real-time adaptive learning predicts the elasticity to be the Kalman gain from a statistical inference problem. Solving recursively backwards

$$\hat{E}_t\pi_{t+1} = g\sum_{j=0}^{\infty}(1-g)^j\,\pi_{t-j}$$

reveals expectations are a distributed lag of current and past inflation rates, with geometrically declining weights, so that lower values of g increase dependence on past inflation.

What are the implications of adaptive expectations for inflation dynamics? We take the Phillips curve, essentially Cagan's model of hyperinflation,

$$\pi_t \;=\; \kappa x_t + \beta \hat{E}_t \pi_{t+1} \tag{2}$$

as a primitive structural equation for analysis. Inflation depends on the current output gap and next-period expected inflation. The parameter β measures the importance of inflation expectations for current inflation, which together with how expectations are formed, determines the dynamic stability of inflation. Substitution for expectations in (2) gives

$$\pi_t = \frac{\beta g}{1-\beta g} \sum_{j=1}^{\infty} (1-g)^j \, \pi_{t-j} + \frac{\kappa}{1-\beta g} x_t.$$

For a given path of the output gap, inflation inherits the distributed lag from expectations.

The learning parameter g regulates the volatility and persistence of inflation. Substituting (2) into the adaptive expectations rule (1) gives

$$\hat{E}_t \pi_{t+1} = \frac{1-g}{1-g\beta} \hat{E}_{t-1} \pi_t + \frac{g\kappa}{1-g\beta} x_t.$$

For $g = 0$ expectations are constant and inflation only reflects movements in the output gap. As the elasticity of expectations rise, expectations become progressively less anchored to past expectations and increasingly more sensitive to current inflation. As a result, the volatility of expectations rises. And for $\beta \to 1$ expectations and inflation become less stable, displaying near-random-walk behavior.[2]

Forecast errors

$$\pi_t - \hat{E}_{t-1} \pi_t = -\frac{g(1-\beta)}{(1-\beta g)(1-g)} \sum_{j=1}^{\infty} (1-g)^j \, \pi_{t-j} + \frac{\kappa}{1-\beta g} x_t$$

also depend on past inflation. That forecast errors are predictable and correlated over time ultimately led to adaptive expectations falling out of favour. Indeed, Muth (1961) argued that "information is scarce, and the economic system generally does not waste it." Because adaptive expectations led to predictable forecast errors available information is wasted. His hypothesis for expectations was based on the principle that predictions of future economic outcomes should be consistent with the relevant economic theory. This idea

[2] For $\beta > 1$ the system becomes unstable, with self sustaining hyper-inflations and deflations.

was to become known as the rational expectations hypothesis and was forcefully advanced by Thomas Sargent and Robert Lucas in the early 1970s.

2.2 Rational Expectations: 1970s

The rational expectations hypothesis states that expectations are given by the mathematical expectations implied by the economic theory. That is $\hat{E}_t \pi_{t+1} = E_t \pi_{t+1}$ where the operator E_t denotes mathematical expectations and the subscript t now implicitly includes the information available when making predictions. In this simple model the information set includes the history of the endogenous inflation rate and the exogenous output gaps up to and including the current period.

This expression defines a first-order linear rational expectations difference equation for inflation. To determine inflation today requires an expectation of next-period inflation. But next-period inflation in turn depends on expectations of inflation two periods hence—and so on into the indefinite future. For this reason, the economic behavior of such models attracted the label "forward-looking". To formulate plans today requires anticipating economic conditions into the future. Mathematically this logic is represented by recursive forward substitution of equation (2) to give

$$\pi_t = \lim_{T \to \infty} E_t \beta^{T-t} \pi_T + \kappa E_t \sum_{T=t}^{\infty} \beta^{T-t} x_T.$$

Current inflation depends on expected discounted long-run inflation plus current and all future expected output gaps.

If $0 < \beta < 1$ the model has a unique bounded solution of the form

$$\pi_t = \frac{\kappa}{1 - \beta\rho} x_t = \frac{\kappa\rho}{1 - \beta\rho} x_{t-1} + \frac{\kappa}{1 - \beta\rho} \varepsilon_t \tag{3}$$

assuming the output gap is a first-order process with ρ the autoregressive coefficient and ϵ_t an *i.i.d.* mean-zero Gaussian innovation.[3] Inflation inherits the statistical properties of the output gap. It follows that

$$E_t \pi_{t+1} = \frac{\kappa\rho}{1 - \beta\rho} x_t.$$

Compared with adaptive expectations, the persistence of expected inflation is tied directly to the persistence of the exogenously given output gap. The parameter β only affects volatility

[3] The requirement of a bounded equilibrium ensures $\lim_{T \to \infty} E_t \beta^{T-t} \pi_T = 0$.

because it regulates the rate at which firms discount future output gaps. Forecast errors

$$\pi_t - E_{t-1}\pi_t = \frac{\kappa}{1-\beta\rho}\varepsilon_t,$$

are both unforecastable and uncorrelated over time. Price setters make efficient use of available information.

While very much accepted in modern macroeconomics, at the time the rational expectations hypothesis raised a number of questions. One concern related to the question of how people would come to hold rational expectations? After all, the rational expectations hypothesis is not a behavioral hypothesis. Rather it simply asserts that people form expectations consistent with the relevant economic theory. How then do people get a handle on complex economic interactions that underlie market outcomes and importantly coordinate with others on this equilibrium. Another related concern was that many rational expectations models predicted a multiplicity of equilibria. Which equilibrium will be reflected in actual economic outcomes? Both considerations gave impetus to further research on the behavioral foundations of expectations formation. In part to resolve indeterminacies of existing theory, but in important other part, to grant plausibility to the rational expectations hypothesis itself.

Lucas (1976) raises further questions of this kind, despite being advanced and frequently cited as a terminal critique of adaptive expectations. Lucas argued that economic models using adaptive expectations were useless for policy analysis. Because actual real-world expectations would surely adjust to incorporate knowledge of the consequences of announced policy changes, macroeconomic models based on historical patterns in data would be misleading. According to this view, a policy change that alters the persistence of the output gap should immediately imply an identical change in the persistence of aggregate inflation. But considerable ground must be covered in moving from the notion of incorporating knowledge of a new policy when forming expectations, to understanding fully the resulting general equilibrium implications in a complex market economy.

Later Lucas (1986) concedes as much, writing

> I think of economics as studying decision rules that are steady states of some adaptive process, decision rules that are found to work over a range of situations and hence are no longer revised appreciably as more experience accumulates.

The introduction of Sargent's Arne Ryde Lectures emphasizes this point arguing that macroeconomics has "given us theories of dynamics that have their best chance of applying when people are in recurrent situations that they have experienced often before." Economists armed with such theory can provide advice about "how to expect a system to operate after it has fully adjusted to a new and coherent set of rules and expectations, but with virtually no theory of the transition itself" going on to say "we might have prejudices and anecdotes .... but no empirically confirmed formal theories" of transitions. We are left then to make judgments about when situations and policies are to be considered familiar and when they are not.

Against this intellectual backdrop, and the concerns about the plausibility and uniqueness of rational expectations equilibrium, was the adaptive learning literature born. Critical of rational expectations, Friedman (1979) articulates the intellectual challenge:

> If a model depends crucially on people's ability to make predictions which possess the error orthogonality property [of rational expectations] ..., it must explain how they come to have the necessary knowledge for this task.

In Benjamin Friedman's view, much of the confusion about rational expectations—and therefore, in our view, alternative conceptions of expectations formation—stems from "failure to distinguish" the efficient use of information from the "specific assumption[s] identifying the available information set".[4] The latter endows people with the knowledge to infer general equilibrium consequences of market economies. While most economists would not dispute efficiency there are various reasons to doubt the knowledge assumptions required for rational expectations. Like Sargent, Friedman was concerned that the underlying behavior of the economy is subject to change because of "new institutional or regulatory or macroeconomic policy" arrangements; that data were subject to measurement problems; and that agents had "finite and imperfect" memory. People grapple with these complexities by using models that are good approximations of the true model over some range of economic outcomes. Indeed, having scrutinized the required knowledge assumptions for rational expectations Roman Frydman concludes "the formulation of individual forecasting behavior in the rational expectations literature appears to be *ad hoc*" (Frydman, 1982, p. 653).[5]

Of course, rational expectations researchers sought their own approach to questions of economic transition. Regime switching models assume change across regimes is represented by some Markov process. This permitted study of structural change and the consequences of different policies. But as argued by Frydman and Phelps (2013), the progress embodied in this research is more apparent than real. In their words: "because they fully specify both the

[4] Where efficiency here is used in the Rawlsian sense.

[5] On a practical note, it is unclear that the Lucas critique ever gained much traction in policy circles. For example, in a summary of the state of affairs in the late 1970s appearing in his recent history of monetary and fiscal policy in the United States, Alan Blinder writes "Academic economists were madly in love with rational expectations, barely recognizing that the term really meant "model consistent" expectations, no matter how silly the model was" (Blinder, 2022, p. 103) Elsewhere he argues the expectational effects of policy announcements are quantitatively unimportant.

process governing change and the post-change representation of outcomes in advance, these models share a key property with their time-invariant counterparts: they describe change with an overarching probability distribution in which the set of outcomes and their associated probabilities are fully predetermined in all time periods—past, present and future." People therefore have complete knowledge of the stochastic processes describing equilibrium outcomes within each regime and the likelihood of each regime occurring. The criticism which regime switching models were intended to address telescope the problem to a grander scale.

2.3 Adaptive Learning: 1980s

Two strands of research emerged out of this work effort exploring departures from full information rational expectations. The first were models of rational learning. Decisionmakers have the correct model of the economy and make model consistent forecasts—that is, they know the pricing functions determined by equilibrium—but face information frictions about certain parameters or the state of the economy.[6] The second were models of adaptive learning. Decisionmakers act like econometricians when constructing forecasts relevant to their decision problems. In contrast to rational learning, the mapping between observed state variables and general equilibrium outcomes is unknown—that is, they don't know the pricing functions characterizing equilibrium outcomes. As new data becomes available they revise their forecasting model using standard statistical tools, such as ordinary least squares. The dependence on realized data gave rise to the label "real-time learning". The remainder of the essay is devoted to this second strand of research.

Bray (1982), Bray and Savin (1986), Fourgeaud, Gourieroux, and Pradel (1986) and others explored the following environment. Observing that rational expectations equilibrium in Cagan's model predicts inflation is a linear function of the exogenous output gap disturbances, assume that decisionmakers forecast future inflation using a regression of past inflation on past output gaps. For an outside observer attempting to uncover the dynamics of rational expectations equilibrium, this model is well specified. However, for a decisionmaker inside the economic model, the model is misspecified. This is because outside of rational expectations equilibrium, inflation depends on the estimates of the decisionmaker's forecasting model, making inflation a time-varying parameter model. Their collective actions change the data-generating process, making the constant coefficient model misspecified. Marcet and Sargent (1989b) were later to call the property of continual interplay between beliefs and the data generating process "self-referentiality".

[6] See for example, Frydman (1982), Brav and Heaton (2015), Bray and Kreps (1987) and Angeletos and Lian (2023) for a survey of the literature.

Formally, decisionmakers adopt the econometric model

$$\pi_t = a_t + b_t x_{t-1} + u_t \tag{4}$$

where u_t is a regression error term. This is the decisionmaker's subjective beliefs about the economic environment. The beliefs nest the rational expectations equilibrium (3). Given estimates of the parameters $\{a_t, b_t\}$ a forecast of future inflation can be constructed as

$$\hat{E}_t \pi_{t+1} = a_t + b_t x_t$$

which assumes people understand the stochastic process for the output gap. When constructing forecasts of future outcomes, decisionmakers assume that the parameters of the forecasting model are constants, not to be revised; even though the statistical model that is used to capture historical patterns in data treats them as uncertain and systematically revised as new data become available. Sargent (1993, 1999) and Kreps (1998) calls this approach to formulating decisions "anticipated utility".

Evaluating expectations in the Phillips curve gives

$$\pi_t = \beta a_t + (\kappa + \beta b_t)\, x_t \tag{5}$$

which is the true data-generating process for inflation. We will call this the objective beliefs, because they would be the correct beliefs for an outside observer of the model who understands the structural determinants of aggregate supply and expectations. The misspecification of the subjective forecasting model is immediate. Decisions based on the belief that the parameters $\{a_t, b_t\}$ are constants induce time-varying parameters $\{\beta a_t, (\kappa + \beta b_t)\,\rho\}$ through the learning process. This mapping from subject to objective beliefs is called the T-map and is a central object in the modern analysis of adaptive learning models.

Following this early literature agents estimate model parameters using ordinary least squares and the available history of data. In the applications in Bray (1982) and Bray and Savin (1986) decisionmakers understood the constant intercept in the above regression to be equal to zero. The regression then took the form

$$b_t = \left(\sum_{j=1}^{t-1} x_{j-1} x_{j-1} \right)^{-1} \sum_{j=1}^{t-1} x_{j-1} \pi_j \tag{6}$$

which can be written recursively as

$$b_t = b_{t-1} + t^{-1} R_t^{-1} x_{t-2} \left(\pi_{t-1} - b_{t-1} x_{t-2} \right) \quad (7)$$

$$R_t = R_{t-1} + t^{-1} \left(x_{t-2} x_{t-2} - R_{t-1} \right). \quad (8)$$

Ordinary least squares gives rise to an elasticity of expectations that is time-varying, and declining over time.[7] Together the equations (5), (7) and (8) completely characterize model dynamics. The central question was under what conditions does the parameter estimate b_t converge to the rational expectations equilibrium $\bar{b}$? That is, under what conditions are the subjective beliefs (4) a consistent estimator, delivering convergence to rational expectations equilibrium?

The contribution of Bray (1982) and Bray and Savin (1986) was to show that for maintained parameter values, there was convergence with probability one. This class of result was important, not only because it gave some justification of rational expectations as the natural outcome of a learning process. It also provided confidence that small expectational errors, relative to rational expectations equilibrium, need not be destabilizing. Because dynamics under least-squares learning are stable, errors would be eliminated over time.

The 1980s closed with a series of papers from Thomas Sargent and Albert Marcet that unified and considerably extended the class of stability results for models of adaptive learning. Marcet and Sargent (1989b) demonstrated that a broad class of a econometric learning procedures were examples of stochastic approximation algorithms. And as such, the convergence theorems of Ljung (1977) apply. Under regularity conditions, the stability of the adaptive learning models are given by an associated ordinary differential equation

$$\frac{\partial \phi}{\partial \tau} = T(\phi) - \phi$$

where ϕ represents subjective beliefs and $T(\phi)$ the objective beliefs implied by the true data generating process.[8] Rational expectations equilibria are stationary points of this difference equation and are the only possible limit points of convergence. Local stability then requires the eigenvalues of the Jacobian $T'(\phi) - I$ to have negative real parts when evaluated at the rational expectations equilibrium of interest.

[7] Here we further refine Hick's definition of the elasticity of expectations to relate to the parameters of the forecast function, as opposed to the forecasts themselves.

[8] See Marcet and Sargent (1989a) and Evans and Honkapohja (2001) for a detailed discussion of the required regularity conditions.

In the Bray and Savin example above, the associated ordinary differential equation is

$$\frac{\partial}{\partial \tau}\begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} \beta a \\ \kappa + \beta \rho b \end{pmatrix} - \begin{pmatrix} a \\ b \end{pmatrix}.$$

The dynamics of the intercept a and the "slope" coefficient b are independent. The two eigenvalues are then $\beta - 1$ and $\rho\beta - 1$, satisfying the requirements for local stability of rational expectations equilibrium. The intuition for these results is illuminated by some simple calculations. Lagging (5) and substituting into (7) gives

$$\begin{aligned} b_t &= b_{t-1} + t^{-1} R_t^{-1} x_{t-2} \left(\left(\kappa + \beta b_{t-1} \right) x_{t-1} - b_{t-1} x_{t-2} \right) \\ &= b_{t-1} + t^{-1} R_t^{-1} x_{t-2} \left(x_{t-2} \left(T \left(b_{t-1} \right) - b_{t-1} \right) + \left(\kappa + \beta b_{t-1} \right) \varepsilon_{t-1} \right) \end{aligned}$$

where we have assumed the intercept is known to be equal to zero for simplicity. For large t the effect of new information ε_t for the estimate b_t is small. The dynamics then are governed by the discrepancy $T(b_{t-1}) - b_{t-1}$, the difference between subjective and objective beliefs. The stochastic approximation literature shows that for large t the dynamics of updating are arbitrarily well approximated by the associated ordinary differential equation. When $T(b_{t-1}) \approx b_{t-1}$ then $b_t \approx b_{t-1}$—beliefs are no longer revised and the elasticity of expectations is zero.

2.4 Adaptive Learning: 1990s

The 1990s witnessed further application of stochastic approximation algorithms to refine and extend results on expectations stability. By considering a richer class of economic models and implied stochastic processes, researchers uncovered a broader set of convergence and non-convergence results. In linear and non-linear models with multiplicity of equilibrium and sunspots, adaptive learning was found to permit convergence to a range of equilibria, even cycles, leaving doubt that behavioral theories of expectations formation could justify the coordination of beliefs on one particular equilibrium. Research also showed that learning might never predict convergence to rational expectations equilibrium but nonetheless imply beliefs remain close to rational expectations values, with the average trajectory of beliefs being given by the associated ordinary differential equation. Results also characterized departures of beliefs from these mean dynamics, so-called large deviations theory.

Evans and Honkapohja (1994a) offers a nice illustration of this work effort. They provide a comprehensive treatment of multiplicity of equilibria and sunspots in linear rational expectations models. In the context of our inflation model, when $\beta > 1$ there are a continuum of equilibria of the form

$$\pi = \beta^{-1}\pi_{t-1} + \gamma x_t + \nu_t - \beta^{-1}\kappa x_{t-1}, \tag{9}$$

where γ is an arbitrary constant and ν_t an arbitrary statistical process. With additional forward-looking expectations, such as $E_t\pi_{t+2}$ in the Phillips curve, the class of solutions in (9) generalize to include the lags in ν_{t-1} and x_{t-2}. Should people employ statistical models which condition on these lagged variables, beliefs can converge to sunspot equilibria. There is no reason to have confidence that beliefs should coordinate on even minimum state variable equilibria. Woodford (1990) and Evans and Honkapohja (1994b, 1995) explore related ideas, showing that in non-linear models adaptive learning can converge to sunspot equilibria and cycles.

These authors along with Paul Romer also highlight that learning need not provide an effective equilibrium selection criterion in a non-linear growth model. Evans, Honkapohja, and Romer (1998) build a model that, because of monopolistic competition and complementarities in different types of capital goods, leads to multiple perfect foresight equilibria with high and low growth. They show that rational expectations equilibrium with growth cycles exist, and that this equilibrium is locally stable under learning. This paper was the apotheosis of a sequence of related papers. For example, Evans and Honkapohja (1993) considers a model with multiple steady states because of local increasing returns in production. In an innovative analysis of policy management of expectations, they show that government spending and seigniorage policy could eliminate certain equilibria.

In other work, Bullard (1992) argued that perhaps we should not expect convergence to rational expectations equilibrium. If people are alert to the fact that the true model has time-varying parameters they can employ appropriate statistical methods. For example, Bullard considers use of both random walk coefficient and return to normality models to estimate the parameter b_t in the inflation model. Intuitively in the random coefficients model people suspect economic conditions might change, warranting a different parameter. There is no convergence to rational expectations.

In the context of our regression model, modify equations (7) and (8) to have a constant gain, with $g > 0$ replacing t^{-1}. In contrast to recursive least squares, which weights all observation equally, this estimator discounts past observations at the rate g. While constant-gain algorithms do not converge, the estimated parameters are related to the underlying rational expectation equilibrium. Evans and Honkapohja (2001) show that, for small gain parameters, estimates are normally distributed with mean equal to the rational expectations equilibrium and variance proportional to the gain. This and the ability to track structural change, made constant-gain algorithms a popular choice in later empirical and quantitative analysis.

There also emerged an interest in the role of policy to deliver convergence to rational expectations equilibrium. Bookending the decade, two pieces of research standout. At the beginning, Howitt (1992) formally addressed Milton Friedman's critique of interest rate control to implement monetary policy. Friedman argued that even small implementation errors of an interest-rate peg by a central bank, would inevitably lead to accelerating inflation or deflation. For example, given inflation expectations, if nominal interest rates were set too low real interest rates would fall relative to the natural rate, leading to an expansion in economic activity and higher inflation. As people marked up their inflation expectations, further real interest rate declines would lead to even higher demand and accelerating inflation. This dynamic was called a Wicksellian cumulative process.

Interpreting the dynamics of beliefs as a general class of adaptive learning algorithms, Howitt (1992) confirmed Friedman's reasoning that an interest-rate peg would lead to instability: such monetary policy rules are inconsistent with convergence to rational expectations equilibrium. The result is the analogue to the Sargent and Wallace (1975) finding that interest-rate pegs lead to indeterminacy of rational expectations equilibrium. While itself interesting, Howitt went further, demonstrating interest-rate rules that respond to inflation were consistent with convergence to rational expectations equilibrium. In contrast to the interest-rate peg, having nominal interest rates respond to endogenous developments in the right way would ensure real interest rates rose in response to inflation pressure. This generalized insights from McCallum (1983) under rational expectations and represents the first statement of the so-called Taylor principle, that nominal interest rates should move more than proportionately to inflation. Howitt's work anticipated much analysis that was to be developed almost a decade later as part of the new Keynesian agenda and stands as a vital and enduring, though under appreciated, contribution to monetary economics.

At the end of the decade, Thomas Sargent shifted the focus from household and firm learning about market conditions, to a policymaker learning about structural economic relationships. Sargent (1999) considered a government that recursively estimated a Phillips curve model. He showed that the model's mean dynamics drove the government to implement the suboptimal time-consistent inflation rate, a self-confirming equilibrium of the model. But that "escape dynamics" recurrently pushed the system away from this equilibrium towards the optimal time-inconsistent inflation rate. These escapes were given clear economic content, with the government temporarily learning a version of the natural rate hypothesis—with no exploitable trade-off in the Phillips curve, the optimal inflation rate was zero. Sargent initially developed these ideas informally, but they were later given formal characterization by Cho, Williams, and Sargent (2002) who showed that, like the mean dynamics, the escape dynamics satisfied a certain ordinary differential equation.[9]

2.5 Adaptive Learning: 2000s and beyond

Learning researchers increasingly turned their attention to questions relating to macroeconomic policy. In part this reflected the new Keynesian synthesis which afforded a coherent model for the analysis of inflation policy. According to this theory, good policy was predicated on the precise management of expectations—leading naturally to assessments of how sensitive policy advice was to what some argued was an overly tight control of expectations. And in other part, the increasing interest among policy makers themselves in analyses that concerned behavioral theories of expectations formation and their practical consequences.[10]

The new Keynesian model drops the assumption that output is exogenously determined and appends to our Phillips curve two additional equations

$$x_t = E_t x_{t+1} - \sigma(R_t - E_t \pi_{t+1} - r_t) \tag{10}$$

$$R_t = \phi_\pi \pi_t + \phi_x x_t \tag{11}$$

where R_t is the nominal interest rate, the instrument of monetary policy; r_t the natural real rate of interest an exogenous process; $\sigma > 0$ the household's elasticity of intertemporal substitution; and $\phi_\pi, \phi_x > 0$ policy parameters. The model determines the equilibrium paths of output, inflation and nominal interest rates.

The first equation describes how interest rates determine demand, providing a theory of the transmission mechanism of monetary policy. It can be derived as the aggregate implication of household intertemporal allocation of consumption and leisure under rational expectations. The second is a monetary policy rule of the kind proposed by Taylor (1993). If the policy parameters satisfy

$$\kappa(\phi_\pi - 1) + (1 - \beta)\phi_x > 0,$$

then real interest rates will rise in response to an increase in inflation.[11] On the theoretical front, seminal work by Bullard and Mitra (2002, 2007) extended and confirmed the insights of Howitt to the new Keynesian environment. On the empirical front, Orphanides and Williams (2005), Milani (2007) and Slobodyan and Wouters (2012) jump-started the effort

[9] The central technical results were based on Noah Williams' dissertation research at the University of Chicago and ultimately published as Williams (2019).

[10] See for example the speech Bernanke (2007).

[11] Woodford (2003) coined this property the "Taylor principle".

to bring new Keynesian models with adaptive learning to the data, with implications for policy analysis. For example, Orphanides and Williams (2005) underscored the quantitative importance of effective management of expectations outside rational expectations. Taylor rules chosen to maximize welfare were more aggressive in response to inflation under constant-gain learning than under rational expectations. Limiting drift in expectations under learning was particularly desirable, because shifting inflation expectations worsen the short-run trade-off between inflation and real economic activity. A now massive literature evolved evaluating theoretical and empirical implications of this model—see Eusepi and Preston (2018b) and Evans (2021).[12]

A second important development during this period concerned the evaluation of different theories of belief formation. This included reduced-form and structural empirical models of survey measures of expectations and also laboratory experiments. We later discuss survey expectations at length because they speak directly to a class of models we discuss next. Building on the seminal paper Brock and Hommes (1997), Cars Hommes emerged as a central figure of a prodigious work program using laboratory experiments to understand how people form expectations and whether rational expectations equilibria can be learned. In comprehensive reviews of the literature Hommes (2013), and more recently Bao, Hommes, and Pei (2021), discuss how the economic environment, and in particular the nature of feedback from aggregate expectations to the economy, determines equilibrium stability. This research also documents pervasive heterogeneity in peoples' expectations. Observed forecasts are best described by a process where decisionmakers switch among different forecasting rules in response to their relative performance, as modeled in Brock and Hommes (1997). Hommes (2021) and Branch and McGough (2018) address the implications of heterogeneous adaptive expectations in macroeconomics and finance, including the new Keynesian framework.

3 Optimal decisions with arbitrary beliefs

But this agenda was not without criticism. In modern macroeconomic models rational expectations equilibrium is premised on two stipulations: i) individual optimization; and ii) mutual consistency of beliefs—that is, subjective and objective beliefs coincide. Sargent's (1993) masterful lectures argued for the learning literature to be understood as relaxing the second requirement. Much of the early monetary policy and learning research unwittingly abandoned both.

Preston (2005) showed that studying recursive representations of economic dynamics implied by rational expectations analysis—aggregate Euler equations—was inconsistent with

[12] For textbook rational expectations treatments see the elegant Gali (2008) and the magisterial Woodford (2003).

individual optimal decisionmaking under arbitrary beliefs. To understand this point we need an economic model of decision making. Assume that there is a continuum of monopolistically competitive firms that produce differentiated goods using labor as the only input to production. When setting their price to maximize profits there is a probability α^{T-t} that the firm will not get to change its price in the subsequent $T-t$ periods. The optimal reset price of a firm in period t is

$$p_t(j) = \hat{E}_t^j \sum_{T=t}^{\infty} (\alpha\beta)^{T-t}[(1-\alpha\beta)\hat{w}_T + \alpha\beta\pi_{T+1}], \tag{12}$$

where $\hat{w}_t$ is the real wage, equivalent to marginal costs, and $\hat{E}_t^j$ the expectations of firm j.[13] Because firms care about the average markup during the period for which the price is anticipated to be fixed, the optimal price depends on anticipated marginal costs into the indefinite future. The firm also is concerned about average inflation because of strategic complementarity: prices that are too far from the aggregate level of prices are undesirable, leading to a loss of market share.

If firms hold the same expectations, $\hat{E}_t^j = \hat{E}_t$ for all j, then those firms re-setting prices in period t face the same optimization problem and therefore choose the same price $p_t^* = p_t^*(j)$. Accounting for the temporal effects of staggered price setting, the optimal price and inflation are related by $(1-\alpha)\hat{p}_t^* = \alpha\pi_t$. Combining with the optimal decision rule determines aggregate inflation as

$$\pi_t = \hat{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t}[\kappa x_T + (1-\alpha)\beta\pi_{T+1}]$$

where $\kappa = (1-\alpha\beta)(1-\alpha)/\alpha$ and the output gap satisfies $x_t = \hat{w}_t$.

The above discussed literature proceeded implicitly making rational expectations assump-

[13] Formally this is simplified version of the supply block of the canonical new Keynesian model. See Preston (2005) for details.

tions. To be explicit these calculations are:

$$\begin{aligned}\pi_t &= E_t\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}[\kappa x_T + (1-\alpha)\beta\pi_{T+1}] \\ &= \kappa x_t + (1-\alpha)\beta E_t\pi_{t+1} + E_t E_{t+1}\sum_{T=t+1}^{\infty}(\alpha\beta)^{T-t}[\kappa x_T + (1-\alpha)\beta\pi_{T+1}] \\ &= \kappa x_t + (1-\alpha)\beta E_t\pi_{t+1} + E_t\sum_{T=t+1}^{\infty}(\alpha\beta)^{T-t}[\kappa x_T + (1-\alpha)\beta\pi_{T+1}] \\ &= \kappa x_t + \beta E_t\pi_{t+1} \qquad (13)\end{aligned}$$

where the third equality invokes the law of iterated expectations in the aggregate probability distribution and the fourth equality makes use of the first line shifted one-period ahead.

The difficulty is that under arbitrary beliefs individual firms do not know the aggregate probability laws. If they did, it is difficult to understand how we are not back at a rational expectations equilibrium, which explains the choice of notation. Firms understand their own objectives and constraints. They have their own forecasting model. While it is true that all firms are the same, firms do not know this to be true. Indeed, analysis of a symmetric equilibrium is a matter of convenience for us as analysts of the model. But there is no presumption that firms within the model have such a sophisticated understanding.[14] In practice firms will be different, having different marginal cost structures and different views of the future. The assumption of symmetry nonetheless captures a critical feature of these more general environments: how is aggregate inflation determined by market equilibrium. The approach of this earlier tradition then is to mechanically replace the mathematical expectations operator—in the equation defined by the final equality—with an alternative belief assumption.

An analogous discussion applies to models of aggregate demand that are based on intertemporal theories of household consumption. Standard models predict optimal consumption decision rules of the form

$$c_t(i) = (1-\beta)\beta^{-1}b_{t-1}(i) + \hat{E}_t^i\sum_{T=t}^{\infty}\beta^{T-t}\left[(1-\beta)\hat{w}_T - \sigma\beta\left(R_T - \pi_{T+1}\right)\right] \qquad (14)$$

where $c_t(i)$ is household i's consumption and $b_{t-1}(i)$ holdings of available financial assets, here assumed to be debt in zero net supply. This expression is an example of permanent

[14]The idea that it is implausible for each firm to have full knowledge of the cross-sectional behavior and expectations of all other firms is also central to Frydman's (1982) non-convergence to rational expectations results under rational learning. The presumption that firms would have such knowledge he labeled ad hoc.

income theory. The optimal consumption depends on current asset holdings plus the present discounted value of future wages. Aggregating across the continuum of households, and imposing market-clearing conditions, gives the aggregate demand equation

$$x_t = \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \sigma \left(R_T - \pi_{T+1} - r_T^n \right) \right]. \tag{15}$$

Only under rational expectations is this expression reducible to equation (10).[15]

The two approaches share the same *learning rule* and thus the same methodology to evaluate the convergence properties of learning dynamics. Where they differ is in the assumed *decision rules* used by households and firms to make spending and pricing decisions. Under rational expectations the two decision rules yield the same optimal decisions. However, in general the economics of each framework is fundamentally different.

The following considerations justify the use of the "infinite horizon" framework. First, by construction the approach is consistent with the underlying model microfoundations. This is not the case for decision rules adopted under the Euler equation approach. The optimal decision rule framework therefore retains interpretability being consistent with solving infinite-horizon intertemporal decision problems which depend on long-term beliefs about policy and macroeconomic fundamentals ignored by the Euler equation approach. For example, the Euler equation approach is inconsistent with the standard characterization of the transmission mechanism of monetary policy embedded in new Keynesian models. The microfoundations stipulate that not only the current interest rate, but also the entire future sequence of expected one-period rates affects spending plans today. Furthermore, in the case that assets, such as public debt, are in positive supply, the Euler equation approach fails to account for the wealth effects on consumption demand through the intertemporal budget constraint. Differences such as these engender different conclusions on various dimensions of policy design.

Second, the anticipated utility approach has the useful property that if the econometric model used by agents to produce forecasts is correctly specified—contains those variables appearing in the minimum-state-variable solution under rational expectations—then the resulting individual behavior is *asymptotically* optimal. That is, behavior under the learning algorithm differs from what would be optimal behavior under the true probability laws by an amount that is eventually arbitrarily small. Since the optimal decision rule is a continuous

[15] As an aside, people often describe learning models as "backward looking". This nomenclature, which sought to differentiate "forward-looking" rational expectations from"backward-looking" adaptive expectations modeling, is somewhat inaccurate when applied to modern microfounded models of economic behavior with adaptive learning. Intertemporal decisionmaking is inherently forward looking, independently of the specific theory of expectations formation.

function of the coefficients of the agents' forecasting rule, beliefs that are arbitrarily close to the correct ones imply behavior that is arbitrarily close to being optimal. This is not a property of the Euler Equation approach.[16]

None of this is to say that the Euler equation approach does not represent the optimal decision rules for some other decision problem or even, for that matter, be a better description of reality—rather, only that it has little to do with the microfoundations of modern macroeconomic models they are asserted to represent.[17] Importantly, it matters: the two approaches can give starkly different policy advice and quantitative predictions. Eusepi and Preston (2018b) provides a comprehensive review of the lessons for policy design.

3.1 Anticipated utility and internal rationality

The anticipated utility solution implies agents do not take into account revisions of future estimates when formulating intertemporal optimal plans. This form of 'irrationality' is a common assumption and made for tractability in many applications. Cogley and Sargent (2008) discuss the implications of anticipated utility in a simple example where the fully Bayesian optimal decision can be computed, concluding the anticipated utility solution well-approximates fully optimal decisions. A closely related solution concept was proposed by Adam and Marcet (2011) which defines fully optimal decisions as 'internally rational'. In their approach, the subjective probability distribution for variables outside the control of decisionmakers may differ from the true data generating process but still imply a dynamically consistent set of beliefs and therefore fully optimal decisions.

To clarify, simplify inflation beliefs in equation (2) to

$$\pi_t = a + u_t \tag{16}$$

so that inflation depends only on an unknown constant mean a and u_t denotes an *i.i.d.* innovation with known variance σ_u^2.[18] Firms recursively estimate the inflation mean using Bayesian learning. The unknown mean is treated as a random variable. In the initial period, prior uncertainty about 'a' is

$$a \sim N\left(a_0, \nu_0^{-1}\sigma_u^2\right), \tag{17}$$

[16] Suppose one half of households have positive wealth which in equilibrium is lent to the other half of households whom are otherwise identical. Then the Euler equation approach stipulates when beliefs converge creditors permanently under consume, and debtors over consume, relative to what is optimal.

[17] Though Section 5 makes clear, only the optimal decision approach can explain salient features of survey data on the term structure of expectations.

[18] This can be relaxed. Given priors for this parameter, agents can update their estimate over time. We omit this discussion as it does not play a role in our linear examples.

a Normal distribution with mean a_0 and variance proportional to the volatility of innovations, with $\nu_0 > 0$ scaling the dispersion of initial beliefs. After observing inflation π_t, the posterior distribution is

$$\begin{aligned} a_{t+1} &= a_t + \frac{1}{\nu_t + 1}(\pi_t - a_t) \\ \nu_{t+1} &= \nu_t + 1, \end{aligned}$$

which describes the law of motion of beliefs, where $g_t \equiv \frac{1}{\nu_t+1}$ is a decreasing gain. The estimator corresponds to recursive ordinary least squares described above. This then provides a microfoundation for the adaptive learning approach, at least in this simple case. Decisionmakers correctly forecast the future evolution of their own beliefs at every horizon, as $E_t a_T = a_t$ for $T > t$. Evaluating expectations in the pricing rule (12) delivers the optimal decision, conditional on their dynamically consistent beliefs.

What if firms believe the mean of inflation is shifting over time? One parsimonious way to capture this variation is to assume the mean of inflation, now denoted $\bar{a}_t$, is approximated by a random-walk

$$\bar{a}_{t+1} = \bar{a}_t + \epsilon_t^a \tag{18}$$

where the innovation has variance σ_a^2. The evolution of the optimal mean estimate of $\bar{a}_t$ can be obtained with the Kalman filter

$$a_{t+1} = a_t + g(\pi_t - a_t) \tag{19}$$

where the $g \equiv \sigma_a^2/\sigma_u^2$ defines the Kalman gain, a function of the signal-to-noise ratio. Again this model provides a dynamically consistent system of beliefs. For asset pricing applications see Adam, Marcet, and Nicolini (2016) and Adam, Marcet, and Beutel (2017). For an application to a New Keynesian model see Eusepi, Giannoni, and Preston (2023b). Sargent and Williams (2005) show more general constant-gain algorithms can be approximated by Kalman filter recursions, although their dynamics can be quite different in small samples.

The two learning algorithms introduce free parameters. In the model with a constant inflation mean, the free parameter is the initial belief a_0. In the random walk model the free parameter is the Kalman gain, corresponding to beliefs about the relative volatility of innovations. The latter is harder to pin down, as it is interpreted as an immutable prior about the economic environment. Agents could estimate these parameters over time. However, the recursive updating of these parameters is more complex to model, especially if we retain the assumption of Bayesian updating.[19]

[19] It is generally hard to characterize analytically the evolution of the posterior distribution of such pa-

Another approach is to select such parameters according to an optimality criterion. However, such criteria require some degree of 'external rationality'. For example, bounds to rationality can be measured by studying the wedge between subjective and objective beliefs that occurs in equilibrium. We discuss these in the following section.

3.2 Optimality and external rationality

Once we depart from rational expectations, what is meant by optimality becomes elusive: it depends on the equilibrium being studied and therefore on the preferences, technologies and expectations of others. Indeed, some such as Woodford (1990) argue "there is no unequivocal meaning for the postulate of "rational" behavior outside of an equilibrium". While we may not want to insist that subjective beliefs coincide with objective beliefs in all situations, so that there is some upper bound on rationality, we may nonetheless desire that expectations satisfy some lower bound on rationality—else we become seduced by the "wilderness of bounded rationality" in which anything is possible.

Marcet and Nicolini (2003) provides one possible approach. A lower bound on rationality should satisfy three principles: i) beliefs should be asymptotically optimal and converge to rational expectations equilibrium; ii) during transition dynamics, forecast errors should not be "too large"; and iii) beliefs are optimal in the sense that conditional on all other agents having these beliefs, it is optimal to also hold the same beliefs. To evaluate, let alone satisfy, these three criteria requires some degree of *external rationality*, a restriction on how much subjective and objective beliefs can differ. Like rational expectations, these criteria come from outside the model. For learning to converge to rational expectations people must be endowed with the 'right' model—requiring knowledge of the equilibrium from the outset!

Of course, we might insist that people test their models in real time like statisticians do, using only information available to them: such an approach requires no assumptions about external rationality. For example, Bray and Savin explore whether people can detect correlation in forecast errors using a Durbin-Watson test in transition. However, Marcet and Sargent (1995), Evans, Honkapohja, Sargent, and Williams (2013) and Cho and Kasa (2015, 2017) demonstrate that standard asymptotic statistical results can not be applied to self-referential learning models.

Moving beyond the simple regression model, decisionmakers might have different forecasting models each better suited to certain conditions.[20] For example, ordinary least squares is appropriate in stationary environments but constant-gain algorithms better suited to pe-

rameters.

[20]See Evans and Honkapohja (2009), Eusepi and Preston (2018b) and Evans and McGough (2020) for a more detailed discussion of model selection and their applications.

riods of structural change. How do people choose among competing models? For many researchers, Bayesian model averaging would be a natural basis for discriminating across models. But Evans, Honkapohja, Sargent, and Williams (2013) and Cho and Kasa (2017) show this procedure can place unit probability weight on the constant-gain model, even when the decreasing-gain model spans the true model of the economy. Brock and Hommes (1997) also show that the optimal choices of forecasting models based of past prediction performance can lead to instability in equilibrium.

To handle these difficulties, Marcet and Nicolini (2003) and Cho and Kasa (2015) propose model selection criteria which lead to selection of the forecasting model which spans the truth with probability one. That the optimality of a particular selection criterion depends on the model and equilibrium being considered, is proof that an external rationality criterion is needed to evaluate specific learning algorithms.

In an economy with constant-gain learning, Evans and Honkapohja (1993) show how external rationality can deliver an optimal gain. In this model agents choose their gain optimally, in the sense that they minimize forecast errors, given the data (which depends on the learning gain chosen by other agents). The optimal gain is then part of the equilibrium of the model and it depends on other structural parameters, such as alternative policy rules. Adopting a similar approach in the New Keynesian model, Lansing (2009) shows the optimal learning gain depends on the volatility and persistence of the output gap process. Marcet and Nicolini (2003) and Carvalho, Eusepi, Moench, and Preston (2023a) extend this approach to environments where agents have to choose not only the parameters but also the appropriate forecasting model and show how specific economic policies alter these choices.

4 So what? Three Examples

This section presents three examples which underscore that models of adaptive learning under optimal decisionmaking provides an intellectual framework to address questions for which rational expectations is unable or generates puzzles. The final section adduces empirical evidence supporting this approach.

4.1 Monetary policy

By assumption rational expectations equilibrium ensures people hold beliefs that are consistent with monetary policy strategy, in the short and long run. Such models can not be used to evaluate central bank communications policy; to evaluate how to anchor long-term inflation expectations; to evaluate under what circumstances would inflation expectations become unanchored; and to address what the answers to these questions collectively imply for inflation policy. Adaptive learning models can (see Preston 2006, Eusepi and Preston 2010, 2012 and Carvalho, Eusepi, Moench, and Preston 2023b).

That long-term expectations are tightly controlled in rational expectations policy analysis has strong implications for policy advice. Adaptive learning models, by relaxing the precision with which expectations can be managed, permits evaluating how sensitive is rational expectations policy advice. In general, unstable long-term inflation and nominal interest rate expectations place constraints on the efficacy of stabilization policy. Not only is aggressive short-run aggregate demand management of the kind proscribed by rational expectations analysis infeasible, but the optimal long-run inflation rate will tend to be higher the more unstable are expectations (Eusepi, Giannoni, and Preston 2018, 2023b). And when the nominal interest rate is constrained by the effective lower bound, forward guidance policy is far less effective at stabilizing the macroeconomy. Relative to rational expectations, learning fundamentally alters wealth effects associated with interest rate policy, which confronts central banks with new trade-offs (Eusepi, Gibbs, and Preston 2022). Eusepi and Preston (2018b) provides a comprehensive review of the consequences of adaptive learning for monetary policy.

Models of adaptive learning assist understanding historical episodes in US monetary history. Eusepi, Giannoni, and Preston (2023b) show such models can explain the high level of long-term interest-rate expectations in the late 1970s, despite the steady decline in Fed policy rate observed in the following decade. Just like professional forecasters at the time, decision-akers in the model fail to anticipate the successful Volcker disinflation and the subsequent formal adoption of an inflation target some 30 years later. Carvalho, Eusepi, Moench, and Preston (2023b) explains the stabilization of long-run inflation expectations in the late 1990 as increased inflation predictability, which lead to an increase in confidence about long-run price stability and the subsequent anchoring of expectations (and consequently, the gradual decline in the learning gain).

4.2 Policy uncertainty and regime change

Adaptive learning models resolve the Frydman-Phelps critique of Markov regime-switching models. Because beliefs are adapted to observed patterns in data, learning models obviate the requirement that all policy regimes, including their likelihood and dynamic properties, are known to decisionmakers. People can hold beliefs about the probability distribution of future outcomes that are in principle consistent with the current policy regime, while also being equally consistent with infinitely many other regimes and alowing adaption to regime change by continuously testing their forecasting models. Crucially, the approach permits study of entirely new regimes. It is this property that permits study of communication and anchoring of expectations.

But the implications are broader. For example, from Leeper (1991) we know that for central banks to be able to control inflation requires fiscal policy to be conducted in a very specific way. But this fiscal backing of monetary policy assumes that people expect the present discounted value of taxes to be exactly equal to the real market value of outstanding public debt. Analyses of monetary policy under rational expectations almost always assume this condition holds, giving rise to the prediction that debt has no monetary consequences and that central banks can control inflation without explicit concern of fiscal policy. But if decisionmakers are uncertain about future tax policy, there will be departures from Ricardian equivalence. Fiscal policy will have monetary consequences even when monetary and fiscal policy are consistent with implementing a Ricardian equilibrium under rational expectations (Evans, Honkapohja, and Mitra 2012 and Eusepi and Preston 2012). Fiscal uncertainty breaks the clean separation between Ricardian and non-Ricardian equilibria, with the result that the scale and maturity structure of public debt matter for inflation, complicating stabilization policy. Such models have been used to evaluate fiscal foundations of trend inflation (Eusepi and Preston 2012, 2018a). Models of this kind predict monetary policy is less effective in periods of high government debt levels and high interest rates. Recent economic developments permit evaluating these ideas, with many countries sustaining unprecedented increases in debt in the aftermath of the 2020 Pandemic, and elevated interest rates as a result of high inflation. Whether elevated debt levels will compromise inflation policy going forward is an open question and an opportunity for further research.

This approach to regime uncertainty extends immediately to models of regime change. For example, Marcet and Nicolini (2003) study how hyperinflations can arise from agents abandoning inflation beliefs consistent with a stable nominal anchor. Similarly they show that beliefs can change abruptly under adaptive learning in the presence of policy change when agents have multiple models.

4.3 Puzzles in Asset Pricing

Rational expectations asset pricing theory base delivers a range of ‘puzzles’, failing to capture the observed volatility of returns and the mean, persistence and predictability of excess returns. Adam, Marcet, and Nicolini (2016) and Adam, Marcet, and Beutel (2017) show that a standard Lucas (1978) asset pricing model with learning can match these empirical regularities.[21] Forecasting expected returns using realized returns, investors become optimistic about the stock market when they observe positive returns, consistent with survey-based expectations. Waves of optimism and pessimism drive stock market fluctuations, so that

[21] See also Branch and Evans (2011) for an alternative modeling framework.

observed price volatility far exceeds that of expected stock dividends. The model also successfully generates the positive co-movement between asset prices and return expectations that we observe in the data. Branch and Evans (2010) show that abandoning the common assumption of a representative investor can lead to additional empirical gains. A simple asset pricing model with heterogenous adaptive expectations can match observed regime switching in asset returns and volatility. Adaptive learning is the source of empirical success.

Adaptive learning models with optimal decisions also permit more fruitful analysis of the interaction between asset prices and the macroeconomy (Du, Eusepi, and Preston 2021, Eusepi, Giannoni, and Preston 2023a, Winkler 2020, Adam and Merkel 2019). Economic models commonly used for macroeconomic analysis make stark predictions about asset prices. Macroeconomic variables of interest can be studied without any knowledge of asset prices. And when asset prices are determined, they are functions of those same fundamentals that determine equilibrium in the macroeconomy. These properties have given rise to various puzzles relating to both the excess volatility of asset prices given relatively stable fundamentals, and the large real economic effects that appear to be associated with some asset price movements. In models of adaptive learning macroeconomic outcomes and asset prices are strongly interdependent. This allows addressing issues such as whether central banks should respond to asset prices—see, for example, Winkler (2020)—and how asset market developments affect the transmission mechanism of monetary policy. The limited work that has been done on these questions suggest a promising future research agenda.

5 Empirical Implications and Lessons

The past thirty years witnessed an explosion of empirical work using survey-based measures of expectations to test rationality and to evaluate alternative models of expectation formation. This work complements theoretical developments, providing further insight into the plausibility of rational expectations as well as possible foundations of behavioral theories of beliefs. Three central lessons emerge from this empirical work. First, the data reject rational expectations. Understanding that all models are false, this is perhaps unsurprising and unhelpful—rational expectations afforded many advantages for macroeconomic modeling, particularly in the absence of clear alternative approaches to modeling beliefs. But the remaining lessons identify a compelling path forward. Second, unobserved components models explain the term structure of expectations of inflation, output growth and interest rates exceptionally well. Moreover, models of the joint evolution of these variables perform better than single equation models, suggesting a degree of sophistication in forecasts. Third, long-term expectations exhibit considerable variation that is correlated with short-run forecast errors. These findings are consistent with models of optimal decision making under adaptive learning.

5.1 How rational are survey-based expectations?

Survey data on expectations have been used in empirical work since the early 1970s. Most studies test whether observed expectations are consistent with rational expectations. In its strongest form subjective beliefs must be the same as the true data-generation process. A weaker implication is that forecast revisions should only incorporate new information: forecast errors should be uncorrelated with past available information—see Friedman (1979) and Nordhaus (1987) for discussion. Observable measures of expectations, even subject to measurement error, allow direct tests of this prediction.

Early work emphasized inflation expectations, in part because of data availability, in other part because, as a component of the Phillips curve, they were central to policy debate. Using *consensus* measures of inflation expectations, inflation forecast errors were serially autocorrelated and predictable based on old information—see Roberts (1997) and references therein.[22] Other studies find similar results for an expanded set of macroeconomic variables, including short-term interest rates, GDP growth and unemployment—see Friedman (1979), Croushore (1998), Eusepi and Preston (2011) and Farmer, Nakamura, and Steinsson (2021). While these results are at odds with the assumption of full-information rational expectations, they are not necessarily inconsistent with agents processing information optimally. Incomplete knowledge about the true data-generating process, perhaps because of structural change or imperfect information about model structural parameters, can result in ex-post correlated forecast errors even if agents learn optimally given available information—see Friedman (1980), Caskey (1985), Lewis (1989).

An influential example of such composition effects is Coibion and Gorodnichenko (2012) which shows that in an economy with many rational *individual* forecasters processing idiosyncratic noisy signals about unobserved inflation, *consensus* measures of inflation expectations under-react to new information relative to rational expectations. Finding inefficient consensus forecasts is not necessarily a rejection of rational expectations at the individual level. However, recent research finds deviations from rationality at the level of *individual forecasters*. Bordalo, Gennaioli, Ma, and Shleifer (2020) use data on individual forecasters to show that they generally over-react to new information relative to rational expectations. Such findings are inconsistent with the rational expectations hypothesis.

Other empirical work shows that survey data on expectations are broadly consistent with professionals forecasting as econometricians would. Statistical models that are updated ac-

[22] The term 'consensus' corresponds to the average inflation forecast across survey participants at any given point in time.

cording to optimality criteria, such as minimization of expected forecast errors, can track the evolution of observed expectations very well. Crump, Eusepi, and Moench (2022) shows a model of this class fits closely the entire term structure of forecasts for inflation, output growth and the short-term nominal interest rate for the US. Angeletos, Huo, and Sastry (2020) study the dynamic adjustment of survey-based forecasts to a business cycle shock. Compared to rational expectations, observed expectations initially under-react to the shock, but tend to over-react in later periods. These results are consistent with subjective expectations differing from the objective probability distribution of the variables they forecast, consistent with the predictions of adaptive learning models.

Even if we grant deviations from strict definitions of rationality, it is important to investigate the *size* of these deviations. Survey data permit this, so informing the bounds to rationality that we impose on models of adaptive learning.

5.2 Bounds to rationality

Empirical analysis finds forecasters are quite sophisticated. New Keynesian models predict agents must *jointly* forecast variables such as inflation, the output gap and the policy rate to make their consumption, saving, labor supply and pricing decisions. The data suggest forecasters take some economic linkages into account. Mullineaux (1980), Gramlich (1983) and Caskey (1985), for example, show that survey-based inflation expectations are informed not only by past inflation, as predicted by the simplest model of adaptive expectations. But also by money growth and other macroeconomic variables, consistent with the predictions of macroeconomic models.

These earlier contributions also show that models with time-varying coefficients better capture the relation between observed expectations and forecast inputs—see Mullineaux (1980) and Branch and Evans (2006). Recent literature provides further evidence that professional forecasts are based on models capturing the joint evolution of macroeconomic variables. Andrade and Le Bihan (2013), Andrade, Crump, Eusepi, and Moench (2016), Crump, Eusepi, and Moench (2022), Crump, Eusepi, Moench, and Preston (2022) provide evidence that multivariate models provide a better fit of the joint behavior of output growth, inflation and short-term interest-rate expectations. For example Crump, Eusepi, and Moench (2022) shows that information contained in inflation and output growth expectations improve substantially our ability to explain expectations of short-term interest rates. This is consistent with agents accounting for the central bank's reaction function.

More generally, survey data can be used to directly elicit bounds on external rationality from observed expectations. Models of adaptive learning estimated using survey data allow joint inference of the parameters governing the learning process and the possibly time-varying wedge between subjective and objective beliefs. Eusepi, Giannoni, and Preston (2023b) estimate a new Keynesian model with learning using a wealth of survey expectations for inflation and the policy rate, at both short- and long-term maturities.[23] The model characterizes the wedge between subjective and objective expectations and which economic shocks determine this wedge.

5.3 Long-term expectations.

Macroeconomic models imply agents are forward-looking and therefore that long-run expectations matter for their economic decisions. Standard models with rational expectations assume market participants live in a well-understood stationary environment. Long-term expectations are constant and play no important role in the dynamic behavior of the economy. This assumption flies in the face of empirical evidence from survey forecasts. Long-term survey forecasts from households, firms, professional forecasters and also central banks display considerable variability, as market participants revise their long-run estimates of labor productivity, output growth, inflation and interest rates in response to forecast errors—see discussion and references in Crump, Eusepi, Moench, and Preston (2022).

Forecasters disagree sharply about long-term outcomes, for some variables like the short-term interest rate even more than in the short run.[24] These observations provide additional evidence against the full-information rational expectations assumption. Market participants are aware of a changing economic environment and appear to track regime changes by responding to new information. Recent work shows that the entire term structure of forecasts, from the short run to the very long run can be captured by a simple model with unobserved trend and cycle components.[25] Forecasters update their long-term forecasts in response to short-term forecast errors, consistent with the type of statistical learning described in this survey. Morover, the relationship between forecast errors and long-run expectations revisions changes over time. The learning gain—the elasticity of expectations—is time-varying and state dependent.[26]

6 Conclusion

This paper traces the history of models of adaptive learning. From modest beginnings in the adaptive expectations literature, models of adaptive learning emerged as part of the rational

[23] See Milani (2023) for a recent survey on dynamic general equilibrium models estimated using survey data on expectations.

[24] Andrade, Crump, Eusepi, and Moench (2016).

[25] See Crump, Eusepi, and Moench (2022) and Crump, Eusepi, Moench, and Preston (2022).

[26] See Carvalho, Eusepi, Moench, and Preston (2023a) and also Cieslak (2018) for a similar result using short-horizon forecasts and Piazzesi, Salamao, and Schneider (2015) for expected bond yields.

expectations revolution to address questions relating to indeterminacy and the plausibility of rational expectations equilibrium. While perhaps failing to deliver a compelling theory of coordination and equilibrium selection for rational expectations, such models have provided invaluable insight about what it means to be rational and how expectations are formed in practice. Moreover, the approach permits addressing questions that rational expectations analysis can not. Importantly, recent empirical work using survey data adduces evidence that supports a models of optimal intertemporal decision making with adaptive learning.

References

Adam, K., and A. Marcet (2011): "Internal Rationality, Imperfect Knowledge and Asset Prices," *Journal of Economic Theory*, 146, 1224–1252.

Adam, K., A. Marcet, and J. Beutel (2017): "Stock Price Booms and Expected Capital Gains," *American Economic Review*, 107(8), 2352–2408.

Adam, K., A. Marcet, and J. P. Nicolini (2016): "Stock Market Volatility and Learning," *Journal of Finance*, 71(1), 33–82.

Adam, K., and S. Merkel (2019): "Stock Price Cycles and Business Cycles," Crc tr 224 discussion paper series, University of Bonn and University of Mannheim, Germany.

Andrade, P., R. K. Crump, S. Eusepi, and E. Moench (2016): "Fundamental Disagreement," *Journal of Monetary Economics*, 83(1), 106–128.

Andrade, P., and H. Le Bihan (2013): "Inattentive Professional Forecasters," *Journal of Monetary Economics*, 60(8), 967–82.

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2020): "Imperfect Macroeconomic Expectations: Evidence and Theory," in *NBER Macroeconomics Annual 2020, volume 35*, NBER Chapters, pp. 1–86. National Bureau of Economic Research, Inc.

Angeletos, G.-M., and C. Lian (2023): "Dampening general equilibrium: incomplete information and bounded rationality," in *Handbook of Economic Expectations*, ed. by R. Bachmann, G. Topa, and W. van der Klaauw, pp. 613–645. Academic Press.

Bao, T., C. Hommes, and J. Pei (2021): "Expectation formation in finance and macroeconomics: A review of new experimental evidence," *Journal of Behavioral and Experimental Finance*, 32, 100591.

Blinder, A. S. (2022): *A Monetary and Fiscal History of the United States, 1961-2021.* Princeton University Press.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020): "Overreaction in Macroeconomic Expectations," *American Economic Review*, 110(9), 2748–2782.

Bowmaker, S. W. (2013): *The Art and Practice of Economics Research: Lessons from Leading Minds*. Edward Elgar Pub.

Branch, W. A., and G. W. Evans (2006): "A Simple Recursive Forecasting Model," *Economics Letters*, (91), 158–166.

——— (2010): "Asset Return Dynamics and Learning," *The Review of Financial Studies*, 23(4), 1651–1680.

——— (2011): "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, 3(3), 159–91.

Branch, W. A., and B. McGough (2018): "Heterogeneous Expectations and Micro-Foundations in Macroeconomics," in *Handbook of Computational Economics*, ed. by C. Hommes, and B. LeBaron, vol. 4, pp. 3–62. Elsevier.

Brav, A., and J. Heaton (2015): "Competing Theories of Financial Anomalies," *The Review of Financial Studies*, 15(2), 575–606.

Bray, M. (1982): "Learning, Estimation and the Stability of Rational Expectations," *Journal of Economic Theory*, 26, 318–339.

Bray, M., and D. M. Kreps (1987): "Rational Learning and Rational Expectations," in *Arrow and the Ascent of Modern Economic Theory*, ed. by G. R. Feiwel, Palgrave Macmillan Books, chap. 19, pp. 597–625. Palgrave Macmillan.

Bray, M., and N. E. Savin (1986): "Rational Expectations Equilibriua, Learning, and Model Specification," *Econometrica*, 54(5), 1129–1160.

Brock, W. A., and C. H. Hommes (1997): "A Rational Route to Randomness," *Econometrica*, 65(5), 1059–1095.

Bullard, J., and K. Mitra (2002): "Learning About Monetary Policy Rules," *Journal of Monetary Economics*, 49(6), 1105–1129.

——— (2007): "Determinacy, Learnability and Monetary Policy Inertia," *Journal of Money, Credit and Banking*, 39, 1177–1212.

Bullard, J. B. (1992): "Time-varying parameters and nonconvergence to rational expectations under leasst square learning," *Economics Letters*, 40, 159–166.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023a): "Anchored Inflation Expectations," *American Economic Journal: Macroeconomics*, 15(1), 1–47.

——— (2023b): "Anchored Inflation Expectations," *American Economic Journal: Macroeconomics*, 15(1), 1–47.

Cho, I.-K., and K. Kasa (2015): "Learning and Model Validation," *Review of Economic Studies*, 82(1), 45–82.

——— (2017): "Gresham's Law of Model Averaging," *American Economic Review*, 107(11), 3589–3616.

Cho, I.-K., N. Williams, and T. J. Sargent (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69(1), 1–40.

Cieslak, A. (2018): "Short-Rate Expectations and Unexpected Returns in Treasury Bonds," *The Review of Financial Studies*, 31(9), 3265–3306.

Cogley, T., and T. J. Sargent (2008): "Anticipated Utility And Rational Expectations As Approximations Of Bayesian Decision Making," *International Economic Review*, 49(1), 185–221.

Coibion, O., and Y. Gorodnichenko (2012): "What Can Survey Forecasts Tell Us about Information Rigidities?," *Journal of Political Economy*, 120(1), 116–159.

Croushore, D. (1998): "Evaluating Inflation Forecasts," Federal Reserve Bank of Philidelphia, Working Paper No. 98-14.

Crump, R., S. Eusepi, E. Moench, and B. Preston (2022): "The Term Structure of Expectations," in *Handbook of Economic Expectations*, ed. by R. Bachmann, G. Topa, and W. van der Klaauw. Elsevier.

Crump, R. K., S. Eusepi, and E. Moench (2022): "The Term Structure of Expectations and Bond Yields," Staff Report 775, Federal Reserve Bank of New York.

Du, Q., S. Eusepi, and B. Preston (2021): "Non-Rational Beliefs in an Open Economy," *Review of Economic Dynamics*, 41, 174–204.

Eusepi, S., M. Giannoni, and B. Preston (2018): "Some Implications of Learning for Price Stability," *European Economic Review*, 106, 1–20.

——— (2023a): "Distoerted Beliefs, Asset Pricing and Monetary Policy," unpublished, University of Texas, Austin.

——— (2023b): "On The Limits of Monetary Policy," unpublished, University of Texas, Austin.

Eusepi, S., C. Gibbs, and B. Preston (2022): "Monetary Policy Trade-offs at the Zero Lower Bound," manuscript, The University of Melbourne.

Eusepi, S., and B. Preston (2010): "Central Bank Communication and Macroeconomic Stabilization," *American Economic Journal: Macroeconomics*, 2, 235–271.

——— (2011): "Expectations, Learning and Business Cycle Fluctuations," *American Economic Review*, 101(6), 2844–2872.

——— (2012): "Debt, Policy Uncertainty and Expectations Stabilization," *Journal of the European Economics Association*, 10(4), 860–886.

——— (2018a): "Fiscal Foundations of Inflation: Imperfect Knowledge," *American Economic Review*, 108(9), 2551–89.

——— (2018b): "The Science of Monetary Policy: An Imperfect Knowledge Perspective," *Journal of Economic Literature*, 56(1), 3–61.

Evans, G. W. (2021): "Theories of Learning and Economic Policy," *Revue d'économie politique*, 131(3), 583–608.

Evans, G. W., and S. Honkapohja (1993): "Adaptive Forecasts, Hyteresis, and Endogenous Fluctuations," *Economic Review*, pp. 3–13.

——— (1994a): "Learning, Convergence, and Satbility with Multiple Rational Expectations Equilibria," *European Economic Review*, 38, 1071–1098.

——— (1994b): "On the Local Stability of Sunspot Equilibriua under Adaptive Learning Rules," *Journal of Economic Theory*, 64, 142–161.

——— (1995): "Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models," *Econometrica*, 63(1), 195–206.

——— (2001): *Learning and Expectations in Economics*. Princeton, Princeton University Press.

——— (2009): "Learning and Macroeconomics," *Annual Review of Economics*, 1(1), 421–449.

Evans, G. W., S. Honkapohja, and K. Mitra (2012): "Does Ricardian Equivalence Hold When Expectations Are Not Rational?," *Journal of Money, Credit and Banking*, 44(7), 1259–1283.

Evans, G. W., S. Honkapohja, and P. Romer (1998): "Growth Cycles," *American Economic Review*, 88(3), 495–515.

Evans, G. W., S. Honkapohja, T. J. Sargent, and N. Williams (2013): "Bayesian Model Averaging, Learning, and Model Selection," in *Macroeconomics at the Service of Public Policy*. Oxford University Press.

Evans, G. W., and B. McGough (2020): "Adaptive Learning in Macroeconomics," in *Oxford Research Encyclopedia of Economics and Finance*.

Farmer, L., E. Nakamura, and J. Steinsson (2021): "Learning About the Long Run," NBER Working Papers 29495, National Bureau of Economic Research, Inc.

Fourgeaud, C., C. Gourieroux, and J. Pradel (1986): "Learning Procedures and Convergence to Rationality," *Econometrica*, 54(4), 845–868.

Friedman, B. M. (1979): "Optimal Expectations and the Extreme Information Assumptions of 'Rational Expectations'Macromodels," *Journal of Monetary Economics*, 79, 23–41.

Frydman, R. (1982): "Towards an Understanding of Market Processes: Individual Expectations, Learning, and Convergence to Rational Expectations Equilibriu,," *The American Economic Review*, 72(4), 652–668.

Frydman, R., and E. S. Phelps (2013): *Rethinking Expetations: The Way Forward for Macroeconomics*chap. Which Way Forward for Macroeconomics and Policy Analysis? Princeton University Press.

Gali, J. (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction of the New Keynesian Framework*. Princeton University Press.

Hommes, C. (2013): *Behavioral Rationality and Heterogenous Expectations in Complex Economic Systems*. Cambridge University Press.

——— (2021): "Behavioral and Experimental Macroeconomics and Policy Analysis: A Complex Systems Approach," *Journal of Economic Literature*, 59(1), 149–219.

Howitt, P. (1992): "Interest Rate Control and Nonconvergence to Rational Expectations," *Journal of Political Economy*, 100(4), 776–800.

Kaldor, N. (1934): "A Classificatory Note on the Determinateness of Equilibrium," *The Review of Economic Studies*, 1(2), 122–136.

Kreps, D. (1998): "Anticipated Utility and Dynamic Choice," in *Frontiers of Research in Economic Theory*, ed. by D. Jacobs, E. Kalai, and M. Kamien, pp. 242–274. Cambridge: Cambridge University Press.

Lansing, K. (2009): "Time Varying U.S. Inflation Dynamics and the New Keynesian Phillips Curve," *Review of Economic Dynamics*, 12(2), 304–326.

Leeper, E. (1991): "Equilibria Under 'Active' and 'Passive' Monetary and Fiscal Policies," *Journal of Monetary Economics*, 27, 129–147.

Ljung, L. (1977): "Analysis of Recursive Stochastic Algorithms," *IEEE Transactions on automatic control*, pp. 551–557.

Lucas, Robert E, J. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46(6), 1429–45.

Lucas, R. E. (1976): "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series on Public Policy*, 11(1), 19–46.

Marcet, A., and J. P. Nicolini (2003): "Recurrent Hyperinflations and Learning," *American Economic Review*, 93(5), 1476–1498.

Marcet, A., and T. Sargent (1995): *Speed of Convergence of Recursive Least Squares Learning with ARMA Perceptions*Basil Blackwell.

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*, pp. 1306–1322.

——— (1989b): "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, (48), 337–368.

McCallum, B. T. (1983): "On Non-uniqueness in Rational Expectations Models: An Attempt at Perspective," *Journal of Monetary Economics*, 11, 139–168.

Milani, F. (2007): "Expectations, Learning and Macroeconomic Persistence," *Journal of Monetary Economics*, 54, 2065–2082.

——— (2023): "Expectational data in DSGE models," in *Handbook of Economic Expectations*, ed. by R. Bachmann, G. Topa, and W. van der Klaauw, pp. 541–567. Academic Press.

Mullineaux, D. J. (1980): "Inflation Expectations and Money Growth in the United States," *The American Economic Review*, 80(1), 149–161.

Muth, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29(3), 315–335.

Nerlove, M. (1958): "Adaptive Expectations and Cobweb Phenomena," *The Quarterly Journal of Economics*, 72(2), 227–240.

Nordhaus, W. D. (1987): "Forecast Efficiency: Concepts and Applications," *The Review of Economics and Statistics*, 69(4), 667–674.

Orphanides, A., and J. C. Williams (2005): "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in *The Inflation Targeting Debate*, ed. by B. S. Bernanke, and M. Woodford. University of Chicago Press.

Piazzesi, M., J. Salamao, and M. Schneider (2015): "Trend and cycle in bond premia," Working paper, Stanford University.

Preston, B. (2005): "Learning About Monetary Policy Rules when Long-Horizon Expectations Matter," *International Journal of Central Banking*, 1(2), 81–126.

——— (2006): "Adaptive Learning, Forecast-Based Instrument Rules and Monetary Policy," *Journal of Monetary Economics*, 53, 507–535.

Roberts, J. M. (1997): "Is Inflation Sticky?," *Journal of Monetary Economics*, 39, 173–196.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press.

——— (1999): *The Conquest of American Inflation*. Princeton University Press.

Sargent, T. J., and N. Wallace (1975): "Rational Expectations, the Optimal Monetary Instrument, and the Optimal Monetary Supply Rule," *Journal of Political Economy*, 83, 241–254.

Sargent, T. J., and N. Williams (2005): "Impacts of Priors on Convergence and Escape Dynamics," *Review of Economic Dynamics*, 8(2), 360–391.

SLOBODYAN, S., AND R. WOUTERS (2012): "Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models," *American Economic Journal: Macroeconomics*, 4(2), 65–101.

WILLIAMS, N. (2019): "Escape Dynamics in Learning Models," *Review of Economic Studies*, 86, 882–912.

WINKLER, F. (2020): "The role of learning for asset prices and business cycles," *Journal of Monetary Economics*, 114(C), 42–58.

WOODFORD, M. (1990): "Learning to Believe in Sunspots," *Econometrica*, 58(2), 277–307.

——— (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.