The Science of Monetary Policy: An Imperfect Knowledge Perspective*

Stefano Eusepi[†]

Bruce Preston[‡]

February 5, 2016

“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.”

Robert E. Lucas Jr

“Rational expectations imposes two requirements on economic models: individual rationality and mutual consistency of perception of about the environment. [...] I interpret a proposal to build models with ‘boundedly rational’ agents as a call to retreat from the second piece of rational expectations (mutual consistency of perceptions) by expelling rational agents from our model’s environment and replacing them with ‘artificially intelligent’ agents who behave like econometricians.”

Thomas Sargent

“The traditional rational-expectations model of inflation and inflation expectations has been a useful workhorse for thinking about issues of credibility and institutional design, but, to my mind, it is less helpful for thinking about economies in which (1) the structure of the economy is constantly evolving in ways that are imperfectly understood by both the public and policymakers and (2) the policymakers’ objective function is not fully known by private agents. In particular, together with the assumption that the central bank’s objective function is fixed and known to the public, the traditional rational-expectations approach implies that the public has firm knowledge of the long-run equilibrium inflation rate; consequently, their long-run inflation expectations do not vary over time in response to new information.”

Ben S. Bernanke

*The authors thank Jim Bullard, Marc Giannoni and Mike Woodford for on-going discussions about much of the material contained here-in, and also a number of anonymous referees and the editors Janet Currie and Steven Durlauf for constructive feedback. Preston acknowledges research support from the Australian Research Council, under the grant FT130101599. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. The usual caveat applies.

[†]Federal Reserve Bank of New York. E-mail: stefano.eusepi@ny.frb.org.

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

1 Introduction

New Keynesian theory identifies a set of principles central to the design and implementation of monetary policy. These principles rely on the ability of a central bank to manage expectations precisely, with policy prescriptions typically derived under the assumption of perfect information and full rationality. In consequence the prevailing policy regime is credible and correctly understood by market participants. Despite considerable advances in understanding, recent events have engendered a re-evaluation of the theory and practice of monetary policy. The challenging macroeconomic environment bequeathed by the financial crisis has led many to question the efficacy of monetary policy, and, particularly, question whether central banks can influence expectations with as much control as previously thought.

One literature particularly well suited to assessing the legitimacy of these concerns are models of imperfect knowledge, or learning dynamics. The assumption of rational expectations is clearly strong: it seems sensible to understand contexts in which subjective beliefs need not exactly coincide with objective probabilities of the economic model. An appealing property of such models is one can meaningfully define anchored expectations as those beliefs consistent with the monetary policy strategy of the central bank. The possibility of ‘unanchored’ expectations presents a challenge for stabilization policy. More generally, policymakers conduct monetary policy facing substantial uncertainty about the economic environment they operate in. A policymaker committed to achieving sound economic policy objectives might still deliver a suboptimal performance in an environment of imperfect knowledge, no matter how well-intentioned and no matter how sophisticated their modeling frameworks.

The objective of this survey is to review what is understood about the challenges to the New Keynesian paradigm posed by imperfect knowledge and to assess the degree of confidence with which one should hold the basic prescriptions of modern monetary economics. The literature is integrated into a single framework where market participants and policymakers optimize given their subjective beliefs. Beliefs are formed using a-theoretical econometric models which are updated in real time, as new information becomes available. A central property of the learning models considered here is the economic system is “self-referential” in the sense of Marcet and Sargent (1989): beliefs affect the data-generating process, which in turn affect beliefs. As expectation errors are propagated through the economy, they become partially self-fulfilling, opening the door to instability. Desirable policies prevent or limit self-sustaining drift in expectations. Dimensions of failure in this regard constitute important practical challenges confronting actual policy making.

The title “The Science of Monetary Policy: An Imperfect Knowledge Perspective” deliberately recalls the seminal piece by Clarida, Gali, and Gertler (1999) also appearing in the *Journal of Economic Literature*, as it summarizes the key principles to be re-examined through the lens of imperfect knowledge. Like that paper, discussion is organized around a sequence of formally stated results to establish two major themes: i) the basic logic of monetary policy design under rational expectations continues to obtain under imperfect knowledge. In particular, the management of expectations remains central; but ii) imperfect knowledge and learning can limit the set of policies available to central banks, rendering aggregate demand management and inflation control in general more difficult than under ra-

tional expectations. Nonetheless from this synthesis emerges a clear policy recommendation: a framework based on some form of price-level targeting that induces appropriate inertia in interest-rate policy is able to approximate optimal policy and is robust to model uncertainty — either in the form of assumed expectations formation; the assumed structural model; or the ability to observe fundamental disturbances. Importantly, the dependence of monetary policy on expectations must be of a very specific kind. For example, simple rules that respond to a measure of inflation expectations, while often desirable in a rational expectations analysis, can lead to macroeconomic instability.

Worth emphasizing is models of learning dynamics have considerable merit being consistent with basic facts about measured expectations and basic facts about forecasting macroeconomic time series. First, as widely documented, survey measures of expectations exhibit autocorrelated forecast errors, which is inconsistent with the strongest form of rational expectations usually assumed in the literature. For example, Roberts (1995) and Adam and Padula (2011) document this for inflation expectations in the US, while Eusepi and Preston (2011) provide evidence for a range of macroeconomic variables. Second, Crump, Eusepi, and Moench (2015) document the evolution of the entire term structure of survey forecasts of inflation, output growth and the short-term interest rate from all available surveys of US professional forecasters. The study finds substantial drift in medium- and long-term expectations for all these variables. As shown in **Figures 1, 2 and 3**, survey-based long-term expectations drift over time. Importantly, some of the variation in short-term forecasts is incorporated in longer-range projections. In contrast, the baseline New Keynesian model implies near-constant medium-term forecasts — see, for example, Gürkaynak, Sack, and Swanson (2005). This drift in long-term expectations reflects both revised beliefs about the long-term behavior of the economy and beliefs about policy. The central leitmotif of what follows concerns how this drift in beliefs modifies optimal policy prescriptions derived under rational expectations.

In writing this piece, we are aware that there are numerous other surveys in the literature.¹ The intention here is to focus specifically on monetary policy in the New Keynesian framework. Our reading of the learning literature is the analysis of monetary policy is one area of research that has progressed sufficiently far to merit such integration and critical reflection. In so doing we hope to identify emerging differences under learning vis-a-vis rational expectations, and show they are manifestations of a common mechanism. A cost of the approach is it will lack breadth, unavoidably failing to cover interesting work in the learning literature. Such work clearly demands attention, but falls outside the purview of the current discussion.

This survey is organized as follows. Sections 2 and 3 review the basic New Keynesian model and the principles underpinning the science of monetary policy under rational expectations. Section 4 introduces the methodological framework for the analysis of learning dynamics and its implications for policy design. Section 5 investigates the robustness of desirable policies under rational expectations to small expectational errors. Section 6 shows how optimal policy which takes full account of agent learning differs from policy advice

¹See, for example, Evans and Honkapohja (2003a), Evans and Honkapohja (2009a), Evans and Honkapohja (2009b) and Woodford (2013).

under rational expectations. Section 7 considers extensions to the baseline framework: the role of central bank communication for stabilization policy, monetary and fiscal policy interactions, and the consequences for policy design of the zero lower bound on the policy rate. Sections 8 and 9 discuss possible further directions for the literature and conclusions.

2 The Basic New Keynesian Model

This section develops a parsimonious version of the canonical New Keynesian model widely used for monetary policy analysis. A range of assumptions are made for expositional simplicity — for example log utility and linear disutility of labor supply. Results for which these assumptions are not pertinent, or lack generality, will be noted where relevant. Further details on the microfoundations can be found in Woodford (2003a) and Gali (2008).

A continuum of households i on the unit interval maximize utility

$$\hat{E}_t^i \sum_{T=t}^{\infty} \bar{C}_T \beta^{T-t} [\ln c_T(i) - \chi n_T(i)],$$

where $0 < \beta < 1$ and $\chi > 0$, by choice of sequences for consumption, $c_t(i)$, and labor supply, $n_t(i)$, subject to the flow budget constraint

$$c_t(i) + b_t(i) \leq (1 + i_{t-1}) \pi_t^{-1} b_{t-1}(i) + W_t n_t(i) / P_t + \Gamma_t(i) / P_t$$

and the No-Ponzi condition

$$\lim_{T \rightarrow \infty} \hat{E}_t^i \left(\prod_{s=0}^{T-t} (1 + i_{t+s}) \pi_{t+s+1}^{-1} \right)^{-1} b_{T+1}(i) \geq 0.$$

The variable $b_t(i) \equiv B_t(i) / P_t$ denotes real bond holdings (which in equilibrium are in zero net supply), i_t the nominal interest rate, $\pi_t \equiv P_t / P_{t-1}$ the inflation rate, W_t is the hourly wage, $\Gamma_t(i)$ dividends from equity holdings of firms and $\bar{C}_T$ exogenous preference shifter. The operator $\hat{E}_t^i$ denotes subjective expectations, which might differ from rational expectations. The latter is defined by the operator $\mathbb{E}_t$.

A continuum of monopolistically competitive firms maximize profits

$$\hat{E}_t^j \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} [p_t(j) y_T(j) - W_T n_T(j)]$$

by choice of $p_t(j)$ subject to the production technology and demand function $y_T(j) = n_T(j) = (p_t(j) / P_T)^{-\theta_t} Y_T$ for all $T \geq t$, with the elasticity of demand across differentiated goods an exogenous process satisfying $\theta_t > 1$; and exogenous probability $0 < \alpha < 1$ of not being able to reset their price in any subsequent period. When setting prices in period t , firms are assumed to value future streams of income at the marginal value of aggregate income in terms of the marginal value of an additional unit of aggregate income today giving the stochastic discount factor $Q_{t,T} = \beta^{T-t} (P_t Y_t) / (P_T Y_T)$.

In a symmetric equilibrium $c_t(i) = c_t = w_t \equiv W_t/P_t = n_t = Y_t$ for all i , $p_t(j) = p_t(j)$ and $b_t(i) = b_t(j) = 0$ for all i, j . For any beliefs satisfying standard probability laws, to a first-order log-linear approximation in the neighborhood of a zero-inflation steady state, optimal individual consumption and pricing decisions can be expressed as

$$\hat{c}_t(i) = \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} [(1-\beta) \hat{w}_{T+1} - (\hat{i}_T - \hat{\pi}_{T+1} - \beta(\bar{c}_T - \bar{c}_{T+1}))] \quad (1)$$

$$\hat{p}_t(j) = \hat{E}_t^j \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [(1-\alpha\beta)(\hat{w}_T + u_T) + \alpha\beta\hat{\pi}_{T+1}] \quad (2)$$

where for any variable z_t , $\hat{z}_t = \ln(z_t/\bar{z})$ the log-deviation from steady state $\bar{z}$, with the exceptions $\hat{p}_t(j) = \ln(p_t(j)/P_t)$, $\hat{i}_t = \ln[(1+i_t)/(1+\bar{i})]$, $u_t = \ln(\theta_t/\bar{\theta})$ and $\bar{c}_t = \ln(\bar{C}_t/\bar{C})$. With a slight abuse of notation, the caret denoting log deviation from steady state is dropped for the remainder, so long as no confusion results. Aggregating across the continuum of households and firms, and imposing market-clearing conditions, the economy is described by the aggregate demand and supply equations

$$x_t = \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} [(1-\beta)x_{T+1} - (i_T - \pi_{T+1} - r_t^n)] \quad (3)$$

$$\pi_t = \hat{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [\kappa(x_T + u_T) + (1-\alpha)\beta\pi_{T+1}] \quad (4)$$

where the output gap is defined as

$$x_t = y_t - y_t^n = w_t$$

the difference between output and the natural rate of output, the level of output determined by a flexible price economy: here $y_t^n = 0$. The associated natural rate of interest $r_t^n = (\beta\bar{c}_t - \hat{E}_t\bar{c}_{t+1})$ is determined by fluctuations in the propensity to consume, which along with the cost-push shock are first-order autoregressive processes, with eigenvalues $0 < \rho_u, \rho_{r^n} < 1$.

The aggregate demand equation determines the output gap as the discounted expected value of future wages, with the second term capturing variations in the real interest rate, applied in future periods, due to changes in the nominal interest rate and goods price inflation. That expected future dividends are irrelevant to consumption plans, to the first-order, reflects the assumption of an infinite Frisch elasticity of labor supply. The aggregate supply curve determines inflation as the discounted future sequence of marginal costs and the inflation rate. The slope of the Phillips curve is measured by $\kappa = (1-\alpha\beta)(1-\alpha)/\alpha$, while the exogenous process u_t captures variations in desired markups reflecting variations in elasticity of demand θ_t .

The model is closed with assumptions on fiscal and monetary policy. For now we abstract from fiscal policy.² Monetary policy is discussed in detail in subsequent sections. To facilitate

²We initially assume a cashless economy with no government spending or debt.

that discussion note under rational expectations the aggregate demand and supply equations can be given the familiar recursive representation

$$x_t = \mathbb{E}_t x_{t+1} - (i_t - \mathbb{E}_t \pi_{t+1} - r_t^n) \quad (5)$$

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t \pi_{t+1} + u_t. \quad (6)$$

The conditions under which relations (3) and (4) are reducible to these expressions are discussed in section 4.4.

3 Basic Principles: Rational Expectations

3.1 Optimal Policy

The policymaker minimizes the loss function

$$\mathbb{E}_t \sum_{T=t}^{\infty} \beta^{T-t} (\pi_T^2 + \lambda_x x_T^2) \quad (7)$$

where $\lambda_x > 0$ determines the relative weight given to output gap versus inflation stabilization. This welfare-theoretic loss function represents a second-order approximation to household utility.³ Feasible sequences of inflation and the output gap must satisfy the aggregate demand and supply equations (3) and (4). Under rational expectations the optimal commitment solution requires satisfaction of the first-order conditions

$$\frac{\lambda_x}{\kappa} x_t + \Lambda_t + \sum_{j=1}^{\infty} a^j \Lambda_{t-j} = 0 \quad (8)$$

$$\pi_t + \Lambda_t - \frac{(1-\alpha)}{\alpha} \sum_{j=1}^{\infty} a^j \Lambda_{t-j} = 0 \quad (9)$$

for all t where Λ_t is the Lagrange multiplier attached to the aggregate demand equation.⁴ The presence of the lagged Lagrange multipliers reflects the well-known property of history dependence, with the infinite sequence of lags reflecting the dependence of optimal decisions on expectations into the indefinite future.⁵ Combining first-order conditions provides

$$\pi_t = -\frac{\lambda_x}{\kappa} (x_t - x_{t-1}). \quad (10)$$

³The weight on the output gap which reflects household preferences $\lambda_x = \kappa/\bar{\theta}$. In the sequel we allow central bank preferences to deviate from households along this dimension.

⁴Attention is confined to optimality from the “timeless perspective” meaning the Central Bank, when contemplating a new commitment, sets policy as if it had in fact committed to this same policy at an earlier date, infinitely far in the past. See Woodford (2003, chap 7) for a detailed discussion.

⁵Because rational expectations equilibrium is typically represented by recursive Euler equations there is only one lagged multiplier in the standard formulation of this problem. The two policy problems are equivalent, however, under the assumption of rational expectations on the part of households and firms.

This consolidated first-order condition — called a target criterion by Giannoni and Woodford (2002, 2010) and Woodford (2003) — requires the central bank to ensure inflation is proportional to the change in the output gap. The lagged output gap captures the history dependence of the optimal commitment solution, a summary statistic of the infinite sequence of Lagrange multipliers in (8) and (9). Under the optimal commitment policy the state-contingent paths for inflation, output and nominal interest rates satisfy

$$\pi_t = (1 - \mu) \frac{\lambda_x}{\kappa} x_{t-1} + \frac{\mu}{1 - \beta\mu\rho_u} u_t \quad (11)$$

$$x_t = \mu x_{t-1} - \frac{\kappa}{\lambda_x} \cdot \frac{\mu}{1 - \beta\mu\rho_u} u_t \quad (12)$$

and

$$i_t = \frac{\lambda_x - \kappa}{\kappa} (1 - \mu) \mu \cdot x_{t-1} + \frac{\lambda_x - \kappa}{\lambda_x} \cdot \frac{\mu(\mu + \rho_u - 1)}{1 - \beta\mu\rho_u} \cdot u_t + r_t^n \quad (13)$$

where μ is the model's only eigenvalue lying in the unit circle.

The target criterion (10) can be recast in terms of the price level as

$$p_t = \bar{k} - \frac{\lambda_x}{\kappa} x_t \quad (14)$$

for some constant $\bar{k}$. Relation (10) is a first-order difference equation for which (14) is the complementary solution. The state-contingent evolutions in each case involve the same equilibrium responses to shocks that occur in each period, and the same long-run average values of all variables.⁶ To keep matters simple set the constant, $\bar{k}$, equal to zero.⁷ Using (14) in (11), (12) and (13) gives the equilibrium paths of prices, the output gap and nominal interest rates for all t .

An important property of the optimal commitment solution is the independence of inflation and the output gap from variations in the natural rate of interest. Only cost-push shocks represent a fundamental trade-off: optimal policy completely stabilizes inflation and the output gap. This property, which is manifest in substantially more complex models used for policy evaluation — see, for example, Justiniano, Primiceri, and Tambalotti (2013) — has been called “Divine Coincidence” — see Blanchard and Galí (2007) and Clarida, Gali, and Gertler (1999).

Complete stabilization of disturbances to the natural rate of interest remains a feature of optimal policy under discretion. Taking expectations as given in the above policy problem implies $\Lambda_{t-j} = 0$ for all $j > 0$ providing a targeting rule of the form

$$\pi_t = -\frac{\lambda_x}{\kappa} x_t. \quad (15)$$

⁶They differ in a transitory, deterministic component of the solution that will depend on initial conditions — that is, on the value of $\bar{k}$. For one particular choice of $\bar{k}$ that depends on the lagged price level at the time a new commitment is being contemplated the solutions will be identical.

⁷This is not important for the learning analysis being a normalization.

Solving with (4) under rational expectations gives the optimal state-contingent paths for inflation and output gap as

$$\begin{aligned}\pi_t &= \frac{\lambda_x}{\kappa^2 + \lambda_x(1 - \beta\rho_u)} u_t \\ x_t &= -\frac{\kappa}{\kappa^2 + \lambda_x(1 - \beta\rho_u)} u_t\end{aligned}\tag{16}$$

which are again independent of the natural rate of interest. Two differences emerge relative to the optimal commitment solution. First, there is no history dependence. Second, the optimal state-contingent response to disturbances differs, a property referred to as stabilization bias.

3.2 Policy Implementation

Sargent and Wallace (1975) and McCallum (1981, 1983) demonstrated policy rules of the form (13), which only respond to exogenous or predetermined variables — sometimes referred to as fundamentals-based policy rules — may imply multiple equilibria. While they are consistent with the optimal equilibrium, they may be equally consistent with many less desirable equilibria. Desirable policies should deliver a unique bounded rational expectations equilibrium. This can be achieved by appropriate conditioning of nominal interest rates on contemporaneous endogenous variables.

A prominent and elegant approach, of which actual inflation targeting regimes are an example, characterizes the central bank’s decision procedure in terms of specific targeting rules: such rules specify a relationship between one or more target variables that must be checked each time an interest-rate decision is made. The instrument setting is deemed appropriate if the specified “target criterion” is satisfied. Giannoni and Woodford (2002) show that a central bank which implements policy according to targeting rules of the form (10), (14) and (15) will deliver a locally unique rational expectations equilibrium for all parameter values under maintained assumptions. To determine the instrument setting in any period the central bank requires a completely specified model of the economy to solve for the equilibrium path of the policy instrument. For example, the target criterion (10) implicitly defines the interest-rate rule

$$i_t = -x_t^* + \mathbb{E}_t \sum_{T=t}^{\infty} \beta^{T-t} [(1 - \beta)x_{T+1} - (\beta i_{T+1} - \pi_{T+1} - r_T^n)]\tag{17}$$

where

$$x_t^* = \frac{\lambda_x}{\lambda_x + \kappa^2} x_{t-1} - \frac{\lambda_x}{\lambda_x + \kappa^2} \left[\mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [\kappa\alpha\beta(x_{T+1} + u_T) + (1 - \alpha)\beta\pi_{T+1}] \right]\tag{18}$$

is the level of output gap jointly satisfying (15) and (4) under rational expectations. Policy depends upon current and projected paths of demand and cost-push disturbances. This choice of nominal interest rate guarantees the target criterion is satisfied each period.

The target criterion approach possesses a number of important characteristics. First, the specific form of optimal targeting rule does not depend on the exact nature of disturbances judged relevant to the economy at any given time. Appropriate interest-rate policy will depend on perceived current and projected disturbances, but the architecture of the target criterion is invariant to these complications of implementation. For an extensive discussion see Svensson (2003). Second, as emphasized by Clarida, Gali, and Gertler (1999), implicit interest-rate rules of this kind have the property that nominal interest rates move more than proportionately to movements in inflation. This point is most easily seen by employing the usual recursive representation of equilibrium conditions under rational expectations. The interest rule defined by (17) and (18) can be expressed as

$$i_t = -\frac{\lambda_x}{\kappa^2 + \lambda_x}x_{t-1} + \left(1 + \frac{\kappa\beta}{\kappa^2 + \lambda_x}\right)\mathbb{E}_t\pi_{t+1} + \mathbb{E}_tx_{t+1} + r_t^n + \frac{\kappa^2}{\kappa^2 + \lambda_x}u_t. \quad (19)$$

This is a defining feature of optimal policy, both under commitment and under discretion, and a property that Clarida, Gali, and Gertler (2000) find holds for US monetary policy in the post-Volcker policy regime. It has been linked to the period of economic stability up to the Great Recession.

A criticism of the target criterion approach is that it relies on the central bank being able to exploit the true structure of the economy to infer the appropriate interest-rate setting. Not only is this informationally intensive, but a policy that is optimal for one model, might nonetheless perform poorly in closely related models. This motivates interest in specifications of policy which are less informationally demanding, while at the same time performing relatively well across a range of models. Indeed, Levin, Wieland and Williams (1999, 2003) and others have argued that particular classes of simple rules exhibit such robustness, while still being consistent with determinacy of rational expectations equilibrium, and being desirable from a welfare perspective.

A prominent example of this kind of policy is the generalized Taylor rule

$$i_t = \rho i_{t-1} + r_t^n + \phi_\pi \pi_t + \phi_x x_t \quad (20)$$

prescribing interest rates to be adjusted in response to current economic conditions, as captured in the contemporaneous inflation rate, output gap and the natural rate of interest, as well as past developments embodied in the previous-period's interest-rate setting. This rule delivers local determinacy of rational expectations equilibrium if and only if

$$\kappa\phi_\pi + (1 - \beta)\phi_x > 1 - \rho. \quad (21)$$

Woodford (2003a) refers to this condition as the Taylor Principle as it requires nominal interest rates to move more than proportionately in response to inflation variations in the long run. An inertial response to past interest rates implies less aggressive responses to current inflation and output are required for determinacy of equilibrium. Taylor's (1993) prescription did not include interest-rate inertia, though Woodford (2003b) demonstrates history dependence permits such rules to approximate the optimal commitment policy for the model developed above. Rotemberg and Woodford (1999) in a closely related model

show first-difference interest-rate rules, rules with $\rho \rightarrow 1$, also approximate fairly closely the fully optimal policy rule. Furthermore, such rules are robust to model uncertainty, performing well across a range of models.⁸

These results constitute the main principles of policy design under a new Keynesian perspective discussed in Clarida, Gali, and Gertler (1999). For future reference, they are summarized here.

1. Absent cost-push shocks optimal policy under commitment or discretion completely stabilizes disturbances to the natural rate.
2. Optimal policy under commitment or discretion implies the policy rate responds more than proportionally to expected inflation.
3. Optimal commitment policy induces history dependence which improves the short-run stabilization trade-off between inflation and the output gap.
4. Policy implementation can be delivered through an inflation targeting rule of the form (10) or an equivalent price-level targeting rule (14), or by simple policy rules of the form (20).
5. Simple policy rules that satisfy the Taylor principle and incorporate a sufficient degree of interest-rate smoothing approximate fully optimal policy under commitment and are robust to model mis-specification.

4 The Methodological Framework

4.1 Modeling Imperfect Information

The rational expectations paradigm comprises two stipulations: (i) agents optimize given their beliefs about the joint probability distribution for various state variables that are independent of their actions and matter for their payoffs; and (ii) the probabilities that they assign coincide with the predictions of the model. This survey adopts an expository framework which retains (i), while replacing (ii) with the assumption that the joint probabilities are formed using an a-theoretical econometric model. Equations (1) and (2) describe the optimal consumption and pricing decisions for arbitrary subjective beliefs about variables outside the control of households and firms. Agents have a completely specified belief system but do not know the equilibrium mapping between the aggregate state of the economy and the variables, exogenous to their decision problems, that they need to forecast. The predictions of their econometric model need not coincide with the predictions of the theoretical model. This follows a now large literature founded on seminal work by Bray (1982), Bray and Savin (1986), Marcet and Sargent (1989), Grandmont (1998), Sargent (1993), Evans and Honkapohja (1994, 1998) and comprehensively reviewed in Evans and Honkapohja (2001).

⁸See Levin, Wieland and Williams (1999, 2003).

Although easily generalized, most models discussed in this survey embed the following assumptions about agents' information. First, agents are assumed to have identical beliefs, though they do not understand this to be true as they have no knowledge of the tastes and beliefs of other agents. By implication, agents cannot infer aggregate probability laws. Second, the econometric model nests the minimum-state-variable solution to the associated rational expectations problem. Third, market participants can perfectly observe aggregate state variables and have full information about the data-generating process of the exogenous aggregate shocks.

As shown in section 3.1, the optimal commitment equilibrium can be written as a linear function of the variables $\{x_{t-1}, u_t, r_t^n\}$. Agents make use of the forecasting model

$$z_t = \omega_{0,t-1} + \omega_{1,t-1}z_{t-1} + \omega_{2,t-1}r_{t-1}^n + \omega_{3,t-1}u_{t-1} + e_t, \quad (22)$$

where $z_t = (\pi_t, x_t, i_t)'$ and the parameters estimated in period t are given by the set of conformable matrices $\omega_t = (\omega_{0,t}, \dots, \omega_{3,t})'$. Expectations can then be evaluated for any future period $T > t$ as

$$\begin{aligned} \hat{E}_t z_T &= (I_3 - \omega_{1,t-1})^{-1} (I_3 - \omega_{1,t-1}^{T-t}) \omega_{0,t-1} + \omega_{1,t-1}^{T-t} z_t \\ &\quad + (I_3 \rho_{r^n} - \omega_{1,t-1})^{-1} (I_3 \rho_{r^n}^{T-t} - \omega_{1,t-1}^{T-t}) \omega_{2,t-1} \rho_{r^n} r_t^n \\ &\quad + (I_3 \rho_u - \omega_{1,t-1})^{-1} (I_3 \rho_u^{T-t} - \omega_{1,t-1}^{T-t}) \omega_{3,t-1} \rho_u u_t \end{aligned} \quad (23)$$

where I_3 is a (3×3) identity matrix.⁹ Beliefs are updated each period as new data become available according to the class of algorithms

$$\omega_t = \omega_{t-1} + g_t t^{-1} R_t^{-1} y_{t-1} (y_t - \omega'_{t-1} y_{t-1})' \quad (24)$$

$$R_t = R_{t-1} + g_t t^{-1} (y_{t-1} y'_{t-1} - R_{t-1})$$

where $y_t = (1, z_t, r_t^n, u_t)'$. Different assumptions about the variable g_t deliver different gains in the filtering problem. When $g_t = 1$ the updating rule (24) is recursive least squares, which gives equal weight to all observed data. When $g_t = \bar{g}t$ recursive updating is given by a constant-gain algorithm, implying past observations are discounted more heavily. An observation n periods old receives a weight of $(1 - \bar{g})^n$. This survey discusses the consequences of both gain assumptions for stabilization policy.

The forecasting models considered in the learning literature embed a key challenge confronting forecasting in practice. Observed changes in macroeconomic data must be decomposed into high- and low-frequency components. Different assumptions about the gain imply

⁹Belief parameters defining the forecast function in period t are predetermined. This avoids an analytically intractable simultaneity which would otherwise arise from the joint determination of beliefs and equilibrium outcomes. This simultaneity is central to the fixed point problem in rational expectations analysis.

different inferences about the relative importance of these components when new information arrives. For example, a constant gain permits tracking new low-frequency patterns in data, such as a change in the policy regime. In contrast, a decreasing-gain algorithm asymptotically ascribes forecast errors to be solely high-frequency movement in data, while the parameters being estimated are time-invariant. Such inference problems are also quite common in macroeconomic theory and empirical applications, starting from Kydland and Prescott (1982). Recent work demonstrates that statistical models with time-varying drifts are able to explain quite well the behavior of output growth, inflation and nominal interest rates — see, for example, Cogley and Sbordone (2008), Stock and Watson (1989, 2007), Cogley and Sargent (2005a), Laubach and Williams (2003), Cogley, Primiceri, and Sargent (2010) and Gürkaynak, Sack, and Swanson (2005). In addition, such models successfully track the evolution of long- and short-term survey forecasts — see Branch and Evans (2006), Edge, Laubach, and Williams (2007), Kozicki and Tinsley (2012) and Crump, Eusepi, and Moench (2015). The belief structure posited in (22) and (24) is consistent with these features.¹⁰

4.2 On the Nature of Irrationality

While one might be prepared to admit beliefs that are in some way misspecified, one might nonetheless insist decisions account fully for the on-going revision of subjective beliefs.¹¹ However, solving for optimal decisions under learning can be extremely complex, as it requires agents to keep track of anticipated future revisions to beliefs when making current decisions.¹² Most of the literature sidesteps this problem by adopting the ‘anticipated utility’ approach to individual optimization (Kreps (1998), Sargent (1999), Preston (2005) and Woodford (2013)). Given current beliefs, decisions in period t are made assuming beliefs are time invariant over the forecast horizon, even though they are in fact subsequently revised. This is one manifestation of the bounded rationality inherent in learning models.

There are notable exceptions. Cogley and Sargent (2008) consider a simple consumption decision problem with imperfect information. They solve for both the anticipated utility and fully optimal solutions and conclude that model dynamics are close for low degrees of risk aversion. Adam and Marcet (2011) solve for the fully optimal solution in a simple asset pricing model with risk neutral investors, where only one-step-ahead forecasts matter for equilibrium determination. Adam, Beutel, and Marcet (2013) extend that analysis to an asset pricing model with risk averse investors. In this case, the learning problem is conveniently re-interpreted as a filtering problem. The forecasting model is a simplified version of (22) where

$$z_t = \bar{\omega}_{0,t-1} + e_t, \quad (25)$$

with the unobserved drift, $\bar{\omega}_{0,t}$, believed to evolve according to a random-walk

$$\bar{\omega}_{0,t} = \bar{\omega}_{0,t-1} + \eta_t$$

¹⁰For a discussion on the correspondences between the algorithm (24) and the Kalman filter see Bullard (1992), Sargent and Williams (2005) and Berardi and Galimberti (2013).

¹¹Adam and Marcet (2011) refer to this solution concept as ‘internal rationality’.

¹²This issue is discussed in detail in Sargent (1999) and, in particular, Cogley and Sargent (2008).

with innovation η_t . Importantly, agents have immutable priors about the constant variances of the innovations e_t and η_t , implying a constant signal-to-noise ratio. In this case, the updating equation from the Kalman filter is

$$\omega_{0,t} = \omega_{0,t-1} + \bar{g}(z_t - \omega_{0,t-1}) \quad (26)$$

where $\omega_{0,t}$ is the current filtered value of $\bar{\omega}_{0,t}$ and $\bar{g}$ is a function of the perceived signal-to-noise ratio. This implies $\hat{E}_t \omega_{0,T} = \omega_{0,t}$ for $T > t$ so that revisions in agents' estimates are unpredictable.¹³

The survey gives focus to belief structures of the form (25) where agents learn only about the drift parameter. This is for three main reasons. First, it considerably simplifies the exposition when developing analytical results. Second, in the literature to be reviewed the drift term generally imposes the strictest requirements on policy for stability.¹⁴ Third, the drift itself has a clean economic interpretation as representing time variation in the perceived long-run conditional mean of each variable of interest. In the case of inflation it reflects uncertainty about a central bank's inflation target. In the case of income, or equivalently wages, it reflects long-run uncertainty about technological advance and productive capacity.

4.3 Stability Analysis: A Simple Example

Consider a flexible-price version of the model described above without cost-push shocks and an i.i.d. natural-rate process ($\rho_{r^n} = 0$). Assume agents understand the interest rate is set according to the rule $i_t = \phi_\pi \pi_t$. Then inflation is determined by

$$\pi_t = \phi_\pi^{-1} r_t^n - \phi_\pi^{-1} \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} [(\phi_\pi \beta - 1) \pi_{T+1}].$$

Agents have the forecasting model of the form (25)

$$\pi_t = \omega_{0,t-1}^\pi + e_t \quad (27)$$

which nests the stationary rational expectations solution of the model. Substituting for agent forecasts, $\hat{E}_t \hat{\pi}_{T+1} = \omega_{0,t-1}^\pi$, for each $T > t$, gives the true data-generating process

$$\pi_t = \phi_\pi^{-1} r_t^n - \frac{\beta - \phi_\pi^{-1}}{1 - \beta} \omega_{0,t-1}^\pi = \phi_\pi^{-1} r_t^n + T(\omega_{0,t-1}^\pi). \quad (28)$$

Notice first that, with $\omega_{0,t}^\pi = 0$ for all t the equilibrium is the rational expectations solution. This is a feature common to all models discussed in this survey. Second, the true data-generating process displays a time-varying drift. The model is “self-referential” in the sense of Marcet and Sargent (1989): beliefs affect the data-generating process, which in turn affect beliefs. This is a key propagation mechanism of learning models. Third, the true

¹³According to the agent's model, the forecast error, $z_t - \omega_{0,t-1}$, has zero mean and a constant variance.

¹⁴While there is no proof of this claim, in our experience researching in this area, we have not seen a counter example.

data-generating process and the self-referential aspect of learning dynamics are regulated by monetary policy. Inspecting (28), the choice of policy coefficient ϕ_π clearly affects the drift of the inflation process. This link between beliefs and policy defines the central challenge learning poses for policy design.

To clarify the effect of policy on the evolution of the economy requires study of the convergence properties of the learning process. Agents update their estimate of the drift according to

$$\omega_{0,t}^\pi = \omega_{0,t-1}^\pi + g_t t^{-1} \left(\phi_\pi^{-1} r_t^n - \frac{1 - \phi_\pi^{-1}}{1 - \beta} \omega_{0,t-1}^\pi \right) \quad (29)$$

where we have substituted for the true data-generating process of inflation. The methodology to analyze the convergence properties of learning dynamics was developed by Marcet and Sargent (1989) and Evans and Honkapohja (2001). The limiting behavior of beliefs can be described by an ordinary differential equation, reflecting the mapping between the perceived drift $\omega_{0,t}^\pi$ in (27) and the actual drift described in (28). The learning literature refers to the implied dynamics as the ‘mean dynamics’. In compact terms, the ODE is

$$\dot{\omega}_0^\pi = T(\omega_0^\pi) - \omega_0^\pi = -\frac{1 - \phi_\pi^{-1}}{1 - \beta} \omega_0^\pi. \quad (30)$$

When agents use recursive least-squares learning, $g_t = 1$, the local stability of the ODE determines whether the learning process converges point-wise to rational expectations. The rational expectations equilibrium is said to be expectationally stable, or E-Stable, if this differential equation is locally stable in the neighborhood of that equilibrium.¹⁵ Equation (30) shows the ODE is stable if and only if $\phi_\pi > 1$, that is the Taylor principle holds. By responding more than proportionally to inflation the policymaker successfully breaks the self-referential mechanism in the model: an increase in inflation expectations is met by a decline in actual inflation (equation (28)) which in turn implies a downward revision in the estimated drift (equation (29)).

Result 1a. *Interest-rate rules of the form (20) deliver locally stable rational expectations equilibria under learning dynamics if and only if they satisfy the Taylor Principle.*

In the case of constant-gain learning, $g_t = \bar{g}t$, the learning process does not converge point-wise to rational expectations as the estimated drift is constantly revised with new information. For values of $\bar{g}$ that are sufficiently close to zero, the ODE (30) determines the conditions for convergence to an invariant distribution centered around the rational expectations equilibrium — for details see Evans and Honkapohja (2001). The result is important because it establishes the property of expectational stability as fundamental to a proper understanding of both decreasing- and constant-gain belief structures. However, for larger values of the constant-gain parameter, describing situations where agents’ forecasts are further away from rational expectations, convergence requirements can be altered. Evans

¹⁵The connection between conditions for E-stability and local convergence of systems under real time recursive learning is discussed in detail in Evans and Honkapohja (2001).

and Honkapohja (2009c) discuss conditions for stability in this case. Noting that with $g_t = \bar{g}t$ relation (29) is a first-difference equation in beliefs, stability is determined by the eigenvalue

$$1 - \bar{g} \frac{1 - \phi_\pi^{-1}}{1 - \beta}.$$

For a given policy rule a sufficiently high value of $\bar{g}$ ensures this eigenvalue lies outside the unit circle, implying instability. This case, which describes expectations that are further away from the equilibrium under full rationality, is discussed in later sections.

4.4 Anticipated Utility and “Euler Equation” Learning

An alternative to the anticipated utility approach, most clearly articulated by Evans, Honkapohja, and Mitra (2013), asserts decisions rules in which only one-period-ahead expectations matter for spending and pricing decisions. The rational expectations equilibrium conditions (31) and (32) are taken as reduced-form behavioral rules governing decisions of households and firms, with expectations determined by subjective one-step-ahead forecasts to give

$$x_t = \hat{E}_t x_{t+1} - \left(i_t - \hat{E}_t \pi_{t+1} - r_T^n \right) \quad (31)$$

$$\pi_t = \kappa x_t + \beta \hat{E}_t \pi_{t+1} + u_t. \quad (32)$$

Examples include Evans and Honkapohja (2001, chap 10) which studies the standard real business cycle model, Bullard and Mitra (2002) which documents E-Stability conditions in the New Keynesian model above, and Slobodyan and Wouters (2012a) which introduces learning in an estimated medium-scale dynamic stochastic general equilibrium model. Evans and Honkapohja (2013) provide a survey of contributions in this literature.

The methodology is referred to as the “Euler Equation” approach, because equations (31) and (32) are suggestive of being the first-order conditions of household and firm maximization problems.¹⁶ 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, this is generally not true under learning: the two decision rules imply different dynamic behavior and they can deliver different convergence properties — the economics of each framework is fundamentally different.

The survey adopts the anticipated utility framework for expository purposes with the understanding that the Euler Equation approach has contributed greatly to the issues explored and is discussed accordingly. In the authors’ view, the following considerations justify the use of an anticipated utility framework. While neither approach is fully optimal, decision rules derived under anticipated utility are by construction consistent with the underlying model microfoundations. This is not the case for decision rules adopted under the Euler

¹⁶Evans, Honkapohja and Mitra (2013), Preston (2005) and Eusepi and Preston (2011a) discuss more broadly the relative merits and limits of these two approaches. Bullard and Eusepi (2014) discuss the different convergence properties each approach implies in a general class of forward-looking models.

equation approach, which exposes the methodology to some limitations.¹⁷ First, because a rational expectations equilibrium can be represented in many different ways, the Euler equation approach is just one of many possible characterizations of economic decisions. Different equivalent rational expectations representations imply different dynamics under learning — see Cho (2014) and Kurozumi and Zandweghe (2014) for some examples. 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 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.¹⁸

Lastly, the anticipated utility framework 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. Of great importance to the subject matter of this survey, the Euler equation approach is inconsistent with the standard characterization of the transmission mechanism of monetary policy embedded in the New Keynesian framework. The model stipulates 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 which are discussed below.

5 Robustness of REE policy advice to learning

This section employs expectational stability analysis to evaluate the robustness of rational expectations policy advice. Robustness is interpreted in the relatively narrow sense of equilibrium selection. Under rational expectations good policies satisfy the criterion of local determinacy or uniqueness of the equilibrium. Expectational stability analysis restricts this class of policy further by imposing an additional criterion: good policies also ensure rational expectations equilibria are robust to small expectational errors.¹⁹ McCallum (2007)

¹⁷For example, Preston (2005) demonstrates that the optimal decision rules (3) and (4) are reducible to (31) and (32) only by applying rational expectations logic.

¹⁸Suppose 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. Evans and McGough (2014) propose “shadow-price learning” which resolves this difficulty while still retaining the assumption of one-period-ahead forecasts.

¹⁹As attention is restricted to policies that ensure determinacy of equilibrium this survey does not review an extensive literature on sunspots and their learnability. See Bullard (2006) and Evans and Honkapohja

and Ellison and Pearlman (2011) demonstrate that for a certain class of model there is a close relationship between these two criteria for equilibrium selection. Bullard and Eusepi (2014) discuss a more general class of model with decision rules based on anticipated utility and with information delays where this equivalence typically breaks down. The following results underscore that E-Stability is generally a stricter requirement than determinacy for monetary policy design and the choice of monetary rule.

5.1 Taylor Principle and Stability

The Taylor Principle has become a central tenant of monetary policy practice. The following generalizes **Result 1a** to the New Keynesian model for the general class of rule (20).

Result 1b. *If monetary policy is specified by policy rules of the form (20), then rational expectations equilibrium is E-Stable if the Taylor Principle (21) holds.*

The foundations of this result were laid out in Howitt (1992) which shows the Taylor Principle ensures expectations stability in a simple monetary model. Bullard and Mitra (2002) establish this result with no interest-rate inertia in the case of the New Keynesian framework in which household and firm decisions depend upon one-period-ahead expectations. Bullard and Mitra (2007) extend their earlier analysis to inertial interest-rate rules.²⁰ Preston (2005) demonstrates the Bullard and Mitra (2002) results continue to hold in a model of anticipated utility. Despite there being many model modifications that can break this principle, even under rational expectations, it constitutes an important robustness result, given its relevance in the literature.²¹

Two additional points are worth underscoring. First, the fundamentals-based rule, delivered as the special case $\rho = \phi_\pi = \phi_x = 0$, is undesirable under both rational expectations and learning. Evans and Honkapohja (2003b) and Preston (2005) provide results on expectations stability. Taken together, these results underscore the rational expectations insights of Sargent and Wallace (1975) and McCallum (1983), regarding implementation of policy using nominal interest-rate rules, continue to be germane under learning dynamics: good policy must condition on the evolution of endogenous variables. Second, rules that exhibit history dependence promote stability under learning dynamics in so far as weaker responses to contemporaneous inflation and the output gap are permitted. While this observation is an immediate implication of **Result 1b**, it holds more generally. Preston (2008) provides a number of related examples in which nominal interest rates respond to the past price level

(2009a) for a review.

²⁰Analytic results for the case with interest-rate inertia have not been established.

²¹For example, Duffy and Xiao (2011) show that including capital alters many basic results under both learning and rational expectations. Gibbs (2015) shows beliefs based on multiple forecasting models can generate multiple equilibria and require more aggressive policy for stability. Branch, Davig, and McGough (2013) examine the stability properties of the New Keynesian model with regime switching in policy rules. The review also assumes a closed economy: there are open-economy extensions such as Llosa and Tuesta (2008), Llosa and Tuesta (2009), Bullard and Singh (2008) and Bullard and Schaling (2005). We also do not discuss questions of hyperinflation — see Marcet and Nicolini (2003) and Adam, Evans, and Honkapohja (2006).

or the output gap. The intuition for this property stems from the self-referential nature of learning dynamics. History dependence ensures monetary policy responds more aggressively to variations in inflation and output, compared to a rule without inertia. This limits self-fulfilling dynamics and associated unanchoring of expectations. The role of history dependence in promoting expectational stability deserves further study, particularly in models with decision rules based on the anticipated utility framework, where aggregate demand depends on the entire path of future expected interest rates. This survey repeatedly returns to this desirable property.

Taylor rules responding to contemporaneous information have garnered criticism on the ground that central banks do not possess such information in practice. For example, McCallum (1999) argues policy should be operational, responding only to endogenous variables that are observed prior to making an interest-rate decision. Additional practical complications arise from key macroeconomic variables being imperfectly observed because of time lags in data collection; because of frequent revisions in estimates provided by statistical agencies; or because variables like the natural rate of output are not directly observable. A central bank might also not want to respond to contemporaneous information because of lags in the transmission of monetary policy. Such delays in evaluating the effects of policy on its target variables poses the problem of how to monitor policy performance. Because of these considerations Svensson (2003) argues that policy should be formulated in terms of forecast targeting. For example, a central bank should aim to achieve targets of the form

$$\hat{E}_{t-j}\pi_{t+k} + \frac{\lambda_x}{\kappa} \left(\hat{E}_{t-j}x_{t+k} - \hat{E}_{t-j}x_{t+k-1} \right) = 0, \text{ for } j, k \geq 0 \quad (33)$$

an immediate generalization of the optimal commitment target criterion. More generally, researchers have explored policy rules that respond to expectations such as

$$i_t = r_t^n + \phi_\pi \hat{E}_{t-j}\pi_{t+k} + \phi_x \hat{E}_{t-j}x_{t+k} \quad (34)$$

with the understanding that the central bank either responds to measured private-sector forecasts, or constructs forecasts in identical fashion. Under rational expectations, rules of this kind are argued to perform well on several dimensions.²² First, for a wide range of parameters they deliver a locally unique (determinate) equilibrium.²³ Second, they have been shown to be robust to model uncertainty by Batini and Haldane (1999) and Levin, Wieland, and Williams (2003), and have been given empirical support by Clarida, Gali and Gertler (1998, 2000). However, imperfect information models indicate some caution is appropriate.

Result 2. *Policy rules of the form (34) with response coefficients implying determinacy under rational expectations may not be robust to learning.*

²²Modified versions of such rules can be designed to implement (33) at least under rational expectations — Eusepi and Preston (2010).

²³Bernanke and Woodford (1997) show local indeterminacy is possible when the response to inflation is overly aggressive. In the class of model considered here, the magnitude of the policy coefficient required for indeterminacy is outside the range of empirically plausible coefficients.

This result due to Preston (2006) is developed further by Eusepi and Preston (2010) in a related analysis. To give a specific example, consider the simple policy rule: $i_t = r_t^n + \phi_\pi \hat{E}_{t-1} \pi_t$. In the benchmark model in the neighborhood of flexible prices ($\alpha \rightarrow 0$)

$$\phi_\pi > \frac{1}{1 - \beta}$$

is a necessary condition for expectational stability.²⁴ This is a radically different prescription for policy, than mere satisfaction of the Taylor Principle. The discount factor of households is central here because it regulates the effects of shifting expectations on household spending plans, and, through the stochastic discount factor, firm price setting. The greater the patience, the more important are beliefs about future macroeconomic conditions for current economic outcomes — requiring a more aggressive policy to ensure price stability.

Further insight is provided by examining the long-term real interest rate, implicitly defined in the New Keynesian model as

$$r_t^L \equiv i_t + \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} (\beta i_{T+1} - \pi_{T+1}).$$

Agents have the forecasting model of the form (25) for inflation, interest rates and income. Consider a small increase in the perceived drift in inflation ($\omega_{t-1}^\pi > 0$) while initial expectations of income and the policy rate remain at steady state ($\omega_{t-1}^x = \omega_{t-1}^i = 0$). Substituting for the perceived drifts in inflation and interest-rate expectations we obtain

$$\begin{aligned} r_t^L &= \phi_\pi \omega_{t-1}^\pi + \frac{1}{1 - \beta} (\beta \omega_{t-1}^i - \omega_{t-1}^\pi) \\ &= \left(\phi_\pi - \frac{1}{1 - \beta} \right) \omega_{t-1}^\pi \end{aligned}$$

where the second inequality substitutes for the initial beliefs about the interest rate. When $\phi_\pi < 1/(1 - \beta)$ the expected path of the real rate declines, leading to higher aggregate demand, and, through the Phillips curve, an increase in inflation which partially validates the increase in expectations.

A crucial difference to a rational expectations analysis of this class of policy rule is the expected path of interest rates does not adjust immediately in response to the shift in inflation expectations. Because beliefs themselves are sluggish, and monetary policy adjusts with a one-period delay, long-term interest-rates increase only gradually.²⁵ This places much of the burden of stabilization policy on current interest rates, rather than expected future rates — under rational expectations small changes in the current policy rate are efficiently transmitted to changes in long-rates to ensure determinacy of equilibrium. Observe also such rules are likely to generate substantial volatility in the policy rate, which would feedback

²⁴This condition also holds for $j = 0$ and $k = 1$ — see Preston (2006).

²⁵Eusepi and Preston (2010) show the instability problem worsens if the private sector makes consumption and pricing decisions with information delays.

into interest-rate expectations. This self-referentiality is particularly important to later discussion of optimal policy and also the conduct of monetary policy when the zero lower bound on nominal interest rates is relevant.

This kind of instability contrasts with Bullard and Mitra (2002) which finds expectations-based rules (34) are consistent with E-Stability when the Taylor Principle holds.²⁶ Their analysis uses equations (31) and (32) as decision rules — the expected path of the interest rate plays no role for aggregate demand determination.²⁷ Bullard and Eusepi (2014) present examples, both under anticipated utility and Euler equation approach, where policy rules of the form (34) can be unstable under learning while inducing a determinate equilibrium. The key factors that break the equivalence are information delays and the presence of a cost channel of monetary policy.²⁸

Result 2 rests on the assumption that the policy rate responds to forecasts produced using the same reduced-form model adopted by private agents. Evidently the use of such forecasts can be undesirable outside of rational expectations equilibrium. Preston (2006) shows that if a central bank instead forms inflation and output gap forecasts conditional on the correct model of the economy, which includes agents' expectations formation mechanism, then the Taylor principle is restored. Forecasting procedures of this kind internalize the consequences of current interest-rate decisions on household and firm expectations, including beliefs about future interest rates. An implication is even when implementing relatively simple instrument rules, a central bank may nonetheless require a sophisticated structural model to conduct monetary policy effectively. This is a feature of any forecast-based decision procedure. Evaluating expectations using a reduced-form model, such as a vector-autoregression, leads to potential instability. This topic is discussed further in subsequent sections.

5.2 Implementing optimal policy under REE

Modern central banks rarely articulate monetary policy strategy in terms of simple rules. Rather, the objective of policy is specified by a numerical target for inflation.²⁹ Furthermore, considerable effort is undertaken: i) to assess the state of the economy; and ii) to convey information about how the target is to be achieved, specifically the economic model, or logic, guiding the choice of interest rate at any given time. Such communication helps manage expectations, and also helps the public evaluate the credibility of announced policy commitments. These strategies are examples of target criteria. This section evaluates the robustness of target criteria considered optimal from the perspective of rational expectations.

²⁶A corollary of this and the Eusepi and Preston (2010) results is indeterminacy need not imply instability under learning.

²⁷Eusepi (2005) and Ascari, Florio, and Gobbi (2014) provide similar examples of instability in models where only one-period-ahead expectations matter, but interest-rate decisions are made one-period in advance. These models differ from Bullard and Mitra (2002) as expectations of one-period interest rates affect consumption plans — not the realized contemporaneous interest rate.

²⁸Bullard and Mitra (2002) and Evans and McGough (2007) also show that determinacy need not imply E-Stability in policy rules that are operational, responding only to lagged variables instead of forecasts.

²⁹For example, the Bank of England pursues an inflation targeting strategy in which policy is set to ensure forecasted inflation is at 2 percent at a two-year horizon. More recently the Federal Reserve announced an inflation target of 2 percent.

Like the previous discussion of simple rules, robustness is interpreted as the conditions required for expectational errors to be eliminated asymptotically.

Suppose the central bank correctly understands agents' behavior, accurately observes all long-horizon forecasts, and attempts to implement the inflation targeting rule

$$\pi_t = \frac{\lambda_x}{\kappa} (x_t - x_{t-1})$$

which is consistent with the optimal commitment equilibrium under rational expectations. The monetary authority therefore understands the true structural relations of the economy are given by (3) and (4). It can guarantee the target criterion is satisfied by adjusting interest rates according to the following implied reaction function

$$i_t = -x_t^* + \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} [(1 - \beta) x_{T+1} - (\beta i_{T+1} - \pi_{T+1} - r_T^n)] \quad (35)$$

where

$$x_t^* = \frac{\lambda_x}{\lambda_x + \kappa^2} x_{t-1} - \frac{\lambda_x}{\lambda_x + \kappa^2} \left[\hat{E}_t \sum_{T=t}^{\infty} (\alpha \beta)^{T-t} [\kappa \alpha \beta (x_{T+1} - u_T) + (1 - \alpha) \beta \pi_{T+1}] \right]$$

is the level of the output gap that jointly satisfies (4) and (10).

The implied reaction function is an example of an expectations-based instrument rule. The appropriate stance of policy at any point in time depends not only on the past output gap, but also on expectations of income, interest rates, inflation, as well as the anticipated component of natural rate and cost-push shocks. Similar calculations and remarks apply to the target criterion (15) for optimal discretion. The discussion of simple instrument rules would perhaps suggest this approach to be susceptible to instability — this is not the case.

Result 3. *If the monetary authority correctly understands the structural equations (3) and (4), and is charged with implementing the inflation targeting rules (10) or (15), then the economy is expectationally stable under learning dynamics for all parameter values.*

Evans and Honkapohja (2003, 2006) and Preston (2008) contain various stability results for this and related target criteria. An immediate implication is the requirements for expectational stability and determinacy of rational expectations equilibrium are identical.³⁰ This is our second important robustness result. Conditional on beliefs being in a sufficiently small neighborhood of rational expectations equilibrium, the target criterion approach, implemented by the implicit instrument rule (35), eliminates expectational errors and ensures convergence of the optimal commitment equilibrium. Woodford (2007) argues this to be a fundamental strength of the target criterion approach.

An obvious question presents itself: why is the instability inherent in simple forecast-based instrument rules not manifest here? The answer lies in the specific nature of dependence of interest-rate policy on expectations. The target criterion approach has the property

³⁰See also McGough, Rudebusch, and Williams (2005) and Woodford (2005) which discuss the merits of target criteria based on long-term interest rates.

that, conditional on current expectations, an interest rate can always be found that ensures satisfaction of the target criterion. Hence the dependence of policy on expectations is of a very specific kind, and one that ensures the objectives of policy are met. This is not the case for simple instrument rules. This confirms implementation of optimal policy requires a substantial amount of information on the part of the central bank.

5.3 Stability under model mis-specification

Of course, a critical assumption is the central bank has knowledge of the true structural relations. Absent this information, the approach can lend itself to greater instability, providing insight into the practical concern of model uncertainty. Examples are given by Evans and Honkapohja (2006) and Preston (2008). Suppose the central bank mistakenly projects the evolution of the economy using (31) and (32) when in fact the true structural relations are (3) and (4). The target criterion then provides the implicit instrument rule (19), albeit with different expectations operator,

$$i_t = \hat{E}_t x_{t+1} - \frac{\lambda_x}{\lambda_x + \kappa^2} x_{t-1} + \left(\frac{\beta \kappa}{\lambda_x + \kappa^2} + 1 \right) \hat{E}_t \pi_{t+1} + \frac{\kappa}{\lambda_x + \kappa^2} u_t + r_t^n.$$

But this rule generates instability for the same reasons as do rules covered by **Result 2** — even though it ensures unique implementation of the optimal commitment equilibrium under rational expectations, and would similarly ensure expectational stability if the true model under learning were given by relations (31) and (32).

What steps can a central bank take to mitigate instability from model uncertainty? Better outcomes can be achieved by the price-level targeting rule (14). As discussed, under rational expectations the target criteria (10) and (14) are equivalent, in so far as they imply the same state-contingent responses to economic disturbances. The difference between these two rules, in the case of learning dynamics, is the price-level targeting rule specifies a different kind of subsequent behavior when one finds that (because the private sector does not behave as they were projected to do) one has failed to achieve the target criterion precisely. Thus the difference between the two rules is a different commitment as to how one will react to seeing that one has missed one's target. Under price-level targeting, current interest-rate policy does not inherit past mistakes — compare the target criterion (10) which depends on the lagged output gap. Freedom from this dependence eliminates scope for self-fulfilling expectations for the same logic underlying simple instrument rules that do not depend on one-period forecasts. The price-level targeting rule is more robust to learning dynamics and suggests optimal monetary policy might best be implemented by explicit reference to the path of the price level rather than the inflation rate.

To facilitate later discussion of the zero lower bound on nominal interest rates, it's useful to have an explicit characterization of a rule implied by price-level targeting. Suppose the central bank projects the evolution of the economy using (31) and (32) when implementing the price-level targeting rule (14). The implied instrument rule is

$$i_t = \frac{\kappa - \lambda_x}{(\kappa^2 + \lambda_x + \beta \lambda_x)} (p_{t-1} + u_t) + \hat{E}_t x_{t+1} + \frac{\beta \kappa + (\kappa^2 + \lambda_x)}{(\kappa^2 + \lambda_x + \beta \lambda_x)} \hat{E}_t p_{t+1} + r_t^n. \quad (36)$$

Again, this rule uniquely implements the optimal commitment solution, and ensures expectational stability for all maintained parameters if the true economy is described by (31) and (32). However, expectational stability is also ensured for all plausible parameter values in the benchmark model. Price-level targeting (PLT) rules display robustness to model uncertainty.

Result 4. *Implementation of price-level targeting implies E-Stability under different models of the economy.*

6 Optimal Policy Under Learning

The previous section evaluates the robustness of optimal policy under rational expectations to small expectational errors. Emphasis was given to convergence: under what conditions will the economy under learning converge to rational expectations equilibrium. This is a rather narrow evaluation of optimal policy. Equally important is to assess whether the proposed policies perform well during the *transition to rational expectations* or in contexts where beliefs fail to converge to rational expectations equilibrium. For example, because economies are subject to constant structural change, agents might periodically update their forecasting models. A simple approach to modelling concern for structural change is perpetual learning, an example of which are the constant-gain algorithms discussed in section 4. This section discusses optimal policy design from the point of view of a policymaker facing an economy where market participants are constantly updating their model of the economy, or where initial beliefs are not too ‘close’ to rational expectations.

6.1 Optimal Simple Rules

Consider the class of targeting rules

$$\pi_t + \psi^{-1}x_t = 0 \tag{37}$$

where ψ is chosen by the monetary authority to maximize agent welfare. Different choices of ψ implicitly define different preferences over output gap and inflation stabilization. Because any choice of policy within this class will not be consistent with the optimal commitment equilibrium, this class of policy is referred to as a simple rule. Assume the central bank has direct control of the output gap as the instrument of policy and correct knowledge of the economy. Moreover, assume no persistence in the cost-push shock u_t ($\rho_u = 0$). The policymaker minimizes the loss function (7) by choice of ψ subject to the Phillips curve (4) and beliefs. Firms forecast inflation and future marginal costs, which are proportional to the output gap, using a model of the form (25)

$$\begin{bmatrix} \pi_t \\ x_t \end{bmatrix} = \begin{bmatrix} \omega_{t-1}^\pi \\ \omega_{t-1}^x \end{bmatrix} + e_t$$

with beliefs revised according to the constant-gain algorithm (26). Evaluating expectations in the Phillips curve and substituting for the current output gap using the targeting rule (37) gives

$$\pi_t = \frac{\alpha\beta\kappa}{(1+\psi\kappa)(1-\alpha\beta)}\omega_{t-1}^x + \frac{(1-\alpha)\beta}{(1+\psi\kappa)(1-\alpha\beta)}\omega_{t-1}^\pi + \frac{\kappa}{1+\psi\kappa}u_t.$$

Because equilibrium beliefs satisfy $\omega_t^x = -\psi\omega_t^\pi$, only expectations about long-term inflation play an important role.³¹ Taking this restriction into account, model dynamics are summarized by the equations

$$\pi_t = T_\pi(\psi) \cdot \omega_{t-1}^\pi + \frac{\kappa}{1+\psi\kappa}u_t \quad (38)$$

$$\omega_t^\pi = \omega_{t-1}^\pi + \bar{g} \left(\left[(T(\psi) - 1) \omega_{t-1}^\pi + \frac{\kappa}{1+\psi\kappa}u_t \right] \right) \quad (39)$$

where

$$T_\pi(\psi) = \frac{\beta}{(1-\alpha\beta)} \left(\frac{1}{(1+\psi\kappa)} - \alpha \right).$$

The second term on the right hand side of (38) corresponds to the inflation response under rational expectations where $\omega_{t-1}^\pi = 0$ for all t — compare the solution under optimal discretion. The first term captures the drift in realized inflation generated by time-varying beliefs. Orphanides and Williams (2005c) refer to drift in beliefs about inflation as *inflation scares*. Cost-push shocks lead to fluctuations in beliefs, which in turn move inflation, partially validating the initial shift in expectations. The stronger the feedback effects, as measured by the drift coefficient, the more self-confirming is any change in beliefs. In the limiting case where the coefficient approaches unity, the drift becomes a random walk: the perceived and actual models of the economy can be arbitrarily close.

A policy rule that is more aggressive towards inflation (one that assigns less weight to output gap stabilization) weakens the feedback from inflation expectations to actual inflation. For sufficiently low values of ψ the coefficient on long-term inflation expectations is positive, implying strong feedback effects. As ψ increases, and policy becomes more aggressive, these feedback effects decline until they are reversed: an increase in inflation expectations is met by a decline in actual inflation. This is the key insight of papers by Bomfim, Tetlow, von zur Muehlen, and Williams (1997), Orphanides and Williams (2005a, 2005b, 2005c, 2007) and Ferrero (2007). Optimal monetary policy in this class of rule prescribes the central bank to be more conservative than it would be under rational expectations, to respond more aggressively to inflation and inflation expectations.

Figure 4 illustrates this point, plotting the variances of inflation and output gap as a function of the parameter $\tilde{\lambda}_x = \psi^{-1}\kappa$. The figure is produced using standard parameters in the literature: $\alpha = 0.7$; $\beta = 0.99$; σ_u^2 , the variance of the cost-push shock, is set to one. These parameters are kept unchanged in the simulations throughout the survey, unless explicitly stated otherwise. Finally, the preference parameter is set to $\lambda_x = 0.25$, while

³¹ Assume for simplicity that initial beliefs are $\omega_0^\pi = \omega_0^x = 0$.

the constant gain is set to $\bar{g} = 0.15$. The solid gray line corresponds to the model under rational expectations, where volatility is minimized with $\psi^{-1} = \lambda_x/\kappa$. The dashed black line, the model under learning, indicates a targeting rule with less weight on the output gap (roughly half in this example) leads to a better outcome. This can be interpreted as a more conservative central banker with $\tilde{\psi}^{-1} = \tilde{\lambda}_x/\kappa$, where $\tilde{\lambda}_x < \lambda_x$. Here the central bank assigns non-trivial weight to output gap stabilization. As the weight λ_x decreases the differences between rational expectations and learning narrow.

The same results obtain if agents revise their beliefs using a decreasing-gain algorithm. The time-varying drifts eventually vanish as agents learn the rational expectations equilibrium. However, as discussed in Tetlow and von zur Muehlen (2001) and Ferrero (2007), policy plays a key role in determining the speed of convergence of the learning process. A targeting rule with an aggressive response to inflation implies fast convergence to rational expectations. Conversely a targeting rule implying higher weight on output gap stabilization induces slow convergence.³² Ferrero (2007) also shows that the asymptotic variance of the estimated drift is larger for less aggressive policies. In other words, the volatility of beliefs is determined by the strength of feedback effects from learning.

Formally, consider a class of monetary policy rule analogous to the expectations-based rule (19) for optimal discretion:

$$i_t = \gamma_\pi \hat{E}_t \pi_{t+1} + \hat{E}_t x_{t+1} + r_t^n.$$

Orphanides and Williams (2005) and Ferrero (2007) emphasize monetary policy rules that respond to *expected inflation* have powerful stabilizing properties. Using the Euler equation approach Ferrero (2007) characterizes the set of adjusted-learning-speed parameters, γ_π^{ALS} , that are consistent with both asymptotic convergence to optimal discretion policy under rational expectations and a further requirement that learning is suitably quick.³³ Parameter values for γ_π^{ALS} that are larger than the corresponding coefficient in the expectations-based reaction function under optimal discretion imply faster convergence and, for many parameter values, higher welfare.

Result 5a. *Under learning, a stronger response to inflation relative to optimal policy under rational expectations improves the stabilization trade-off.*

Orphanides and Williams (2007, 2008) push this analysis further in an empirical New Keynesian model with habit formation and price indexation where learning is implemented using the Euler Equation approach. Using a simulation-based analysis, two main conclusions emerge which generalize **Result 5a**. First, optimal control policies derived under rational expectations perform poorly when the true model has constant-gain learning. The key reason behind this failure is such policies imply a response to inflation that is too mild to stabilize inflation expectations under learning. Optimal control policies under rational expectations that would arise under a more conservative central banker deliver a better

³²As shown in Marcet and Sargent (1989), if the coefficient on the drift $T_\pi(\cdot)$ is less than 0.5 then the agents' estimate of the drift converges at the same speed as it would in absence of self-referential effects.

³³See also Ascari, Florio, and Gobbi (2014) for a related analysis in models with trend inflation.

performance under learning, consistent with **Figure 4**. Second, simple operational policy rules, which respond to expected or lagged inflation, lagged unemployment gaps and lagged interest rate deliver a consistently good performance under learning. A key feature of such rules is the relatively strong response to inflation and a high degree of interest smoothing — further evidence that interest-rate inertia is desirable under both rational expectations and learning dynamics, and robust to model uncertainty.

6.2 Optimal policy under full control

The previous section explores optimal policy design within a restricted class of policies. The analysis now turns to characterizing the fully optimal policy under learning dynamics. Assume the central bank has rational expectations and has complete knowledge of the true structural relations describing household and firm behavior. Interpret this as a best-case scenario. To the extent that learning dynamics impose constraints on what the central bank can achieve, these difficulties will only be more acute with limited information. Moreover, the nature of these constraints might inform the choice of simple policy rules.

Continue with the thought experiment of a central bank with direct control of the output gap as the instrument of policy. The consequences of modeling the transmission mechanism is taken up in the sequel. The policymaker minimizes the loss function (7) subject to the Phillips curve (4) with i.i.d cost-push shocks and the evolution of beliefs ω_t^π and ω_t^x , updated according to the constant-gain algorithm

$$\begin{bmatrix} \omega_t^\pi \\ \omega_t^x \end{bmatrix} = \begin{bmatrix} \omega_{t-1}^\pi \\ \omega_{t-1}^x \end{bmatrix} + \bar{g} \left(\begin{bmatrix} \pi_t \\ x_t \end{bmatrix} - \begin{bmatrix} \omega_{t-1}^\pi \\ \omega_{t-1}^x \end{bmatrix} \right). \quad (40)$$

In contrast to a rational expectations analysis beliefs are now part of the state space — there is no distinction between commitment and discretion. The loss function continues to be the welfare-theoretic loss implied by a second-order approximation to utility. Imperfect knowledge does not affect these calculations. Analyzing the optimality conditions in an economy in the neighborhood of flexible prices so that $\alpha \rightarrow 0$; and using the fact that the optimal welfare-theoretic weight on the output gap is $\lambda_x = \kappa/\bar{\theta}$ — see Woodford (2003a) — permits derivation of the consolidated first-order condition

$$\pi_t = -\bar{\theta}^{-1} \left(x_t - \bar{g}\beta^2 \mathbb{E}_t \sum_{T=t}^{\infty} [\beta(1-\bar{g})]^{T-t} x_{T+1} \right). \quad (41)$$

This target criterion, first derived by Molnar and Santoro (2013) under the Euler Equation approach, states inflation should be proportional to the output gap, adjusted for the expected discounted value of the sequence of output gaps into the infinite future. The relevant discount is given by the product of the household’s discount factor, β , and one minus the gain.

Several insights are forthcoming. First, for gains in the neighborhood of zero, $\bar{g} \rightarrow 0$, optimal policy is equivalent to optimal discretion under rational expectations — compare equation (15). The intuition is immediate. For gain coefficients that are nearly zero, beliefs

are almost never revised: policy cannot influence them. The policy problem is then formally equivalent to optimal discretion under rational expectations. This logic also holds for the asymptotic behavior of an economy where agents update their beliefs using a time-dependent gain sequence, $g_t = t^{-1}$. Second, for a given gain coefficient, policy will be more aggressive towards inflation relative to optimal discretion. Consider a sequence of positive forecasts for inflation (negative expected output gaps). The target criterion stipulates inflation should be lower today relative to the optimal policy under rational expectations. Policy should limit drift in inflation expectations by keeping inflation close to target, implicitly defined as price stability. This result, consistent with the findings of Ferrero (2007) and Orphanides and Williams (2005a), underscores learning dynamics present an additional trade-off confronting monetary policy. Drifting beliefs constrain what can be achieved by the monetary authority today.

Result 5b. *Optimal policy reveals an additional intertemporal trade-off. Policy is more responsive to inflation to limit future movements in inflation expectations which worsen the short-run trade-off between inflation and the output gap.*

The source of this trade-off resides in the sluggish adjustment of beliefs which induce history dependence in the optimal adjustment of the policy instrument. The solution for the output gap and inflation under optimal policy takes the form

$$\begin{bmatrix} \pi_t \\ x_t \end{bmatrix} = \bar{c}_\pi \omega_{t-1}^\pi + \bar{c}_x \omega_{t-1}^x + \bar{c}_u u_t, \quad (42)$$

where $\bar{c}_\pi$, $\bar{c}_x$ and $\bar{c}_u$ are vectors, functions of the underlying structural parameters. Using (42) to solve for expectations, and iterating (40) recursively backwards permits (41) to be expressed as

$$\pi_t = -\bar{\theta}^{-1} x_t + \sum_{j=1}^{\infty} (1 - \bar{g})^j (\tilde{c}_\pi \pi_{t-j} + \tilde{c}_x x_{t-j}) - \tilde{c}_u u_t,$$

where $\tilde{c}_\pi$, $\tilde{c}_x$ and $\tilde{c}_u$ are again convolutions of model parameters. Regardless of the specific values of the parameters, this formulation emphasizes that optimal policy under learning shares some of the features of optimal policy under commitment in the case of rational expectations. To better see this, **Figure 5** displays impulse responses of inflation and the output gap to a one standard deviation shock to u_t , under different policy configurations.³⁴ The top panels describe impulse responses when the weight on output gap stabilization is consistent with the model's microfoundations ($\lambda_x = \kappa/\bar{\theta} \simeq 0.014$). The black solid line denotes the impulse response with optimal policy under learning, the solid gray line shows the response with the targeting rule (10) under learning. Finally, the dashed black line describes the impulse response under optimal commitment. The two policies under learning imply very similar responses, because the weight on the output gap consistent with microfoundations is fairly small. Similarly to optimal commitment under rational expectations, optimal policy under learning drives inflation below its steady state, although

³⁴Recall the parameters used are the same as for Figure 4.

the size of the decline is smaller. Recall this decline is absent under optimal discretion, which prescribes inflation and output gap to be zero from the second period and beyond.

The response of the output gap is different. The output gap exhibits inertia, but, in contrast with the commitment solution, converges monotonically from above steady state. Lower inflation leads to a downward revision of the *entire path* of inflation expectations in (4), as agents extrapolate a lower inflation mean, increasing the output gap in equilibrium. Under rational expectations inflation expectations decline only in the short-term, leading to a smaller but still negative output gap. When the central bank has a stronger preference for output stabilization (bottom panels), inflation remains positive under the optimal learning policy, but somewhat below inflation implied by the targeting rule that does not fully take expectations formation into account. Overall, optimal policy is closer to the commitment solution than under the simple targeting rule or optimal discretion.³⁵

The gains from policy inertia are taken up further in Gaspar, Smets and Vestin (2007, 2010) and Mele, Molnar, and Santoro (2012). Gaspar, Smets and Vestin consider a model in which only one-period-ahead expectations matter to spending and pricing decisions. In contrast to the above discussion, beliefs include additional state variables and are non-linear, formed by least-squares regression. They show that optimal policy conditional on this more general belief structure continues to be inertial, and continues to be more aggressive towards inflation compared to optimal discretion. Mele, Molnar and Santoro (2013) adopt a similar modeling strategy to find that the optimal commitment equilibrium is not optimal under learning: for decreasing-gain learning algorithms the optimal long-run policy under learning is optimal discretion. As discussed in section 7.5, because the commitment solution yields the highest welfare gain when a zero lower bound is considered, further research should evaluate the robustness of this result in such model environments.

Result 5c. *Optimal policy under learning displays history dependence which results in tighter inflation control. However, it is a different kind of history dependence than what is implied by optimal commitment under rational expectations. In particular, expectations need not display history dependence beyond that implied by the updating rule for beliefs.*

A final property of the Molnar and Santoro optimal policy problem is policy delivers Divine Coincidence — disturbances to the natural rate of interest are completely stabilized. In the anticipated utility environment this property emerges as a special case, but fails to hold more generally. Suppose firm production is subject to an aggregate exogenous productivity shock. Aggregating the optimal pricing rules of firms yields the Phillips curve

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

³⁵These results rely on the central bank having correct knowledge of the agents' expectations formation mechanism. However, Molnar and Santoro (2013) show that these policy prescriptions are robust to policymaker uncertainty about the expectations formation process. In particular, using either a robust control max-min or Bayesian model averaging criterion they show that optimal policy under learning performs well even when private agents have rational expectations.

where a_t denotes technology shocks. In this model the output gap is not the same as the real wage. Formally $x_t = y_t - y_t^n = y_t - a_t = w_t - a_t$. If agents forecast real wages according to

$$\omega_t^w = \omega_{t-1}^w + \bar{g}(x_t + a_t - \omega_{t-1}^w)$$

an additional shock enters in the central bank's constraint set, creating a trade-off: full stabilization of both inflation and output gap is not possible. Because technology shocks affect beliefs about the permanent level of real wages, firm's marginal cost structures shift as households adjust labor supply. Positive technology shocks are stagflationary — see Eusepi, Giannoni, and Preston (2012). The next section reveals this stabilization challenge becomes more acute once one models the transmission mechanism of monetary policy.

6.3 Modeling the Transmission Mechanism

A large literature in monetary economics concerns the transmission mechanism: how do changes in short-term interest rates, the policy instrument under control of the policy-maker, affect economic decisions? The previous sections side-step this question by assuming the central bank has direct control of the output gap. This assumption is now relaxed. In the context of the current model, the transmission mechanism operates through aggregate demand: current and expected future interest rates regulate current spending and savings plans. This section demonstrates that additional trade-offs confront policymakers when account is taken of the transmission mechanism. Importantly, the case for aggressive monetary policy is weakened, and indeed, reversed: relative to optimal discretion, interest rates should adjust by less in response to economic disturbances. As a consequence, Divine Coincidence does not hold — it is no longer possible to fully stabilize inflation and the output gap, even in absence of technology shocks.

To introduce the result, recall when implementing the optimal targeting rule under rational expectations (10), the implied reaction function is given by (35). This rule reveals *policy expectations* play a key role in the ability of the central bank to stabilize the economy. The effectiveness of monetary policy in the New Keynesian model lies in the ability to manage the entire term structure of interest-rate expectations. Under rational expectations, interest-rate beliefs are consistent with the monetary policy strategy of the central bank, permitting tight control of aggregate demand. Under imperfect information, policy expectations need not be consistent with policy strategy, with important consequences for policy design. Reflecting this difficulty, and in contrast to the E-Stability results discussed in section 5, Eusepi, Giannoni, and Preston (2012) show that under constant-gain learning, targeting rules of the form (15) cannot be implemented: reaction functions of the form (35) imply instability under learning for many plausible parameter values — recall from section 4.3 that E-Stability and stability conditions under constant-gain learning need not be the same.³⁶

The key to this instability is the dynamics of policy expectations. In the benchmark model, households and firms must forecast inflation, income and interest rates. Assume

³⁶See Evans and Honkapohja (2009c) for a detailed analysis of this issue.

agents update their estimate of the drifts according to

$$\omega_t = \omega_{t-1} + \bar{g}(z_t - \omega_{t-1}) \quad (43)$$

where $\omega_t = [\omega_t^\pi \ \omega_t^x \ \omega_t^i]'$ and $z_t = [\pi_t \ x_t \ i_t]'$. To keep things simple suppose there are only i.i.d. shocks to the natural rate r_t^n and beliefs are initially consistent with rational expectations equilibrium providing $\omega_{t-1}^\pi = \omega_{t-1}^x = \omega_{t-1}^i = 0$. It is easy to verify that the aggregate demand constraint defines the implicit policy rule

$$i_{t+1} = r_{t+1}^n - \frac{\beta}{1 - \beta} \omega_t^i$$

which is consistent with full stabilization $x_{t+1} = \pi_{t+1} = 0$ in every period t . Substituting into the updating rule for beliefs, ω_t^i , gives

$$\omega_{t+1}^i = \left(1 - \frac{\bar{g}}{1 - \beta}\right) \omega_t^i + \bar{g} r_{t+1}^n$$

which is a first-order difference equation. Sustainable policy requires the dynamics of beliefs to be stationary, implying that the following restriction must hold

$$\bar{g} < 2(1 - \beta).$$

Learning gains $\bar{g}$ violating the above condition lead to explosive interest rates, clearly not a permissible, or at least desirable, feature of optimal policy, if only because the zero lower bound on interest rates obviates such solutions.

Result 6. *In the benchmark model given by (3) and (4), and shocks to the natural interest rate, Divine Coincidence does not hold even in absence of cost-push shocks.*

This result, due to Eusepi, Giannoni, and Preston (2015), stands in direct opposition to the earlier literature. The crucial difference when compared to rational expectations and the Euler equation approach to learning is that the optimal policy problem includes aggregate demand (3) as a constraint, because interest-rate beliefs are also a relevant state variable.

Optimal policy, for plausible gains, stipulates less aggressive responses to economic developments. The transmission mechanism embodies a new intertemporal trade-off confronting optimal policy. The difficulty emerges from drift in interest-rate beliefs, and in consequence, long-term interest rates themselves. In response to a natural-rate disturbance the central bank would like to increase current interest rates. Subsequently, in revising beliefs, agents interpret some of the increase as permanently higher long-term nominal interest rates. This response in the yield curve in turn generates instability in aggregate demand. For higher gain coefficients this difficulty becomes more acute. A given sized disturbance translates into a larger revision to the permanent component of beliefs making it optimal to move interest rates even less.

Figure 6 shows the standard deviations of output gap (dashed lines) and interest rate (solid lines) as a function of the constant gain $\bar{g}$ under the optimal policy. Only i.i.d.

variations in the natural rate, r_t^n , drive economic fluctuations. The figure describes outcomes under the loss function (7), captured by the black lines, and under a loss function that also penalizes volatility in the interest rate (gray lines).³⁷ Recall optimal discretion corresponds to the case where $\bar{g} = 0$. Under the standard loss function a knife edge result obtains: for $\bar{g} < 2(1 - \beta)$ the output gap is fully stabilized even if this induces substantial volatility to the interest rate. For larger values of $\bar{g}$, the policymaker loses the ability to stabilize the output gap. Feasible policy permits limited variation in the policy rate, translating into increasing volatility of the output gap. If the policymaker has some preference for interest-rate stabilization (perhaps reflecting zero lower bound considerations), then the increase in output volatility occurs continuously with the size of the gain. Even relatively small values of the gain lead to considerable output gap volatility.

Figure 7 shows the volatility of inflation and output (weighted with λ_x) as a function of λ_x , for three different optimization problems. The economy is driven by both i.i.d. cost-push and natural-rate shocks. The first policy (gray solid line) is optimal discretion, the second policy (dashed black line) is optimal policy under learning with the central bank having direct control of the output gap. The third policy is the optimal policy subject to the aggregate demand constraint. The figure suggests, even in this very simple modeling framework, the quantitative impact of the aggregate demand constraint on monetary policy effectiveness can be nontrivial. Under full control of the output gap, optimal policy under learning tracks quite closely optimal discretion, at least for sufficiently low values of λ_x .

The consequences of the aggregate demand constraint are not specific to optimal policy. For example, **Figure 8** describes the stability properties under constant-gain learning of a simple Taylor rule, where interest rates respond to inflation and the output gap — an example of ‘robust learning stability’ proposed by Evans and Honkapohja (2009c). Each of the three lines describe the stability frontier in the constant gain and inflation response coefficient space. Parameter regions above a plotted contour indicate local instability of the equilibrium. The black solid line corresponds to $\phi_x = 0.5/4$ (the standard Taylor rule), while the solid and dashed grey lines correspond to output response coefficients which are half and a third of the standard rule respectively. The figure reveals an aggressive stance on inflation, advocated in **Results 5a** and **5b**, might not be feasible. Secondly, a lower response to the output gap enhances stability for ranges of ϕ_π close to empirically plausible estimates.³⁸ These results suggest further investigation of the quantitative implications of drifting interest-rate beliefs has clear merit. In particular, understanding the advantages, if any, of interest-rate inertia, or policies that depend on the lagged price level would be beneficial. Optimal policies within these classes of rule should also be studied.³⁹

³⁷The weight on interest-rate volatility is $\lambda_i = 0.004$. The model parameters are kept at the same values used in the previous simulations.

³⁸Evans and Honkapohja (2009c) obtain a similar result under a forecast-based policy rule.

³⁹For example, Evans and McGough (2007) study the E-stability properties of a class of optimized Taylor-type rules. Constant-gain algorithms are not considered.

6.4 Optimal policy with central bank learning

A defining feature of virtually all interest-rate decision procedures evaluated in this survey is the requirement that a central bank must have a structural model of the economy to implement policy. This is true of forecast-based Taylor rules, the inflation and price-level targeting rules, and also the various conceptions of optimal policy under learning.⁴⁰ All these policies require the central bank to take a stand on the projected evolution of the economy, given the policy framework, when choosing the appropriate interest rate in any given period. This complication distinguishes the policymaker's problem from that of households and firms, which, understanding their own objective and constraints, need only employ an a-theoretical statistical model to forecast variables exogenous to their decision problem. Some examples of the consequences arising from a central bank with a misspecified model have already been given. For example, a central bank that makes interest-rate decisions based on a rational expectations model, when in fact agents are learning. But this kind of misspecification is quiet on the issue of how a central bank might adjust their structural model in the light of poor performance. This section explores the consequences of central bank learning about the structural relations governing aggregate conditions.

Evans and Honkapohja (2003b) consider a central bank which sets optimal discretionary policy using an expectations-based instrument rule, as described in section 5. In contrast to earlier discussion, the policymaker has incomplete information about the true structure of the economy. Specifically, it lacks knowledge of two key structural parameters which govern the transmission mechanism of monetary policy: the interest elasticity of aggregate demand and the slope of the Phillips curve. Because the central bank observes all aggregate data, including expectations, it can recursively estimate the structural parameters, ϕ and κ , using the regressions

$$\begin{aligned} x_t - \hat{E}_t x_{t+1} - r_t^n &= -\phi \left(i_t - \hat{E}_t \pi_{t+1} \right) + e_{x,t} \\ \pi_t - u_t - \beta \hat{E}_t \pi_{t+1} &= \kappa x_t + e_{\pi,t} \end{aligned}$$

where $e_{\pi,t}$ and $e_{x,t}$ are error terms. The central bank's reaction function

$$i_t = \left(1 + \frac{\hat{\kappa}_{t-1} \beta}{(\hat{\kappa}_{t-1}^2 + \lambda_x) \hat{\phi}_{t-1}} \right) \hat{E}_t \pi_{t+1} + \hat{\phi}_{t-1}^{-1} \hat{E}_t x_{t+1} + \frac{\hat{\kappa}_{t-1}}{(\hat{\kappa}_{t-1}^2 + \lambda_x) \hat{\phi}_{t-1}} u_t$$

has time-varying coefficients being functions of the current parameter estimates $\hat{\phi}_{t-1}$ and $\hat{\kappa}_{t-1}$. Estimation errors feedback into the economy through policy decisions, in turn affecting the beliefs of households and firms. However, given that the policymaker can observe private expectations, its model of the economy remains well-specified even during the learning process, unlike the private sector. This assumption limits the impact of central bank learning on the convergence process. And, for this reason, even when both the private sector and the central bank are learning, convergence to rational expectations obtains.

⁴⁰Technically, any decision procedure in which nominal interest rates respond to endogenous contemporaneous information requires the central bank to have a structural model. This includes the standard Taylor rule.

A similar conclusion is reached in Dennis and Ravenna (2008) under different informational assumptions. The policymaker does not observe private-sector expectations and assumes incorrectly that expectations are formed rationally. The two structural parameters are estimated using a general instrumental variable estimator, which imposes the orthogonality restrictions implied by rational expectations. Despite model mis-specification, numerical analysis shows that the economy converges to rational expectations. Policymaker learning can induce higher or lower welfare relative to rational expectations, depending on initial beliefs of both the central bank and the private sector.

These studies mainly focus on asymptotic outcomes under central bank learning. Other contributions extend these ideas to investigate the role of central bank learning in explaining the evolution of inflation in the post-war period, with particular emphasis on the Great inflation and the Volcker disinflation. At the core of these theories is an optimal policy framework in which policy deliberations are made with a potentially misspecified Phillips curve. Starting with the seminal work of Sargent (1999) various papers study how changing beliefs about the short-run inflation-unemployment trade-off can explain the evolution of US inflation in the 1970s and 1980s. The analyses are motivated by the observation that during the 1960s policymakers believed in an exploitable trade-off between unemployment and inflation and choose policy accordingly. The subsequent decade with high inflation and unemployment then induced belief in a version of the natural-rate hypothesis where no exploitable trade off exists, leading to the Volcker disinflation.

To give content to these ideas, introduce a Phillips curve of the form

$$\pi_t = \mu\pi_{t-1} + (1 - \mu)\mathbb{E}_{t-1}\pi_t + \kappa x_t + u_t \quad (44)$$

which can be derived by assuming a constant fraction of goods prices are fully flexible, while the remaining fraction are set one period in advance, on the basis of period $t - 1$ information, along with a fraction $(1 - \mu)$ of firms forming rational expectations, while the remaining fraction sets expectations in purely adaptive form: $\hat{E}_{t-1}\pi_t = \pi_{t-1}$. The output gap is assumed to be imperfectly measured, reflecting the practical challenge of obtaining an estimate of the natural level of output y_t^n . (In the papers discussed here the Phillips curve is expressed in terms of the unemployment gap, but this is not important for our discussion.) The central bank updates its estimate of y_t^n according to

$$\hat{y}_t^n = \hat{y}_{t-1}^n + \bar{g}_y (\hat{y}_t - \hat{y}_{t-1}^n), \quad (45)$$

reflecting the belief that the natural level of output evolves according to a random walk. Orphanides (2001), Orphanides and Williams (2002) and Laubach and Williams (2003) document and discuss the evidence for very noisy estimates of natural rates.

Denote the estimate of the output gap as $\hat{x}_t = \hat{y}_t - \hat{y}_t^n$ in each period t . This is the policy instrument of the central bank. The central bank does not know the true structural Phillips curve. Instead, policy is determined using the empirical model

$$\pi_t = \omega_{0,t} + \omega_{\pi,t-1}\pi_{t-1} + \omega_{x,t-1}\hat{x}_t + \hat{u}_t, \quad (46)$$

where $\hat{u}_t$ denotes the innovation. The model is estimated each period using the recursive least-squares algorithm of the form (24) described in section 4. Finally, the monetary policy

minimizes (7) subject to (46), which results in the target criterion

$$\pi_t + \frac{\lambda_x}{\omega_{x,t-1}} \hat{E}_t (\hat{x}_t - \beta \omega_{\pi,t-1} \hat{x}_{t+1}) = 0. \quad (47)$$

Primiceri (2005) shows this class of model explains fairly well the rise and fall in inflation in post-war US. Three periods are identified. The first, referred to as the period of over optimism, is characterized by the wrong beliefs that the natural level of output was high relative to the truth, coupled with a low estimate of inflation persistence, $\omega_{\pi,t}$. Consistent with these beliefs, optimal policy (47) tolerates rising inflation until the early 1970s. Because of this expansionary policy, during the second period, which spans the 1970s, policymakers revised upward their estimate of inflation persistence. A higher estimate of $\omega_{\pi,t}$ is consistent with a stronger response to inflation under optimal policy: this requires a lower level of the current output gap (given the higher weight on future negative gaps). However, at the same time as updating beliefs to a more persistent inflation process, policymakers revised downward substantially their estimate of the slope of the Phillips curve, $\omega_{x,t}$. This reflects the joint rise of unemployment (a negative output gap) and inflation during the period. Again from (47), a flatter Phillips curve implies a weaker inflation response, which prevented policymakers aggressively fighting inflation, delivering a period of over pessimism. Not until the 1980s did an upward revision of $\omega_{x,t}$ towards its true value lead to a successful disinflation.

In a closely related analysis, Sargent, Williams, and Zha (2006) show a combination of macroeconomic shocks and the central bank's evolving beliefs can explain the rise and fall of inflation in the US during the 1970s and 1980s. A key difference relative to Primiceri is the direction of fit in the estimation of the Phillips curve. In Sargent, Williams, and Zha (2006) the policymaker estimates a new classical version of (44), where the output gap is the dependent variable. A policymaker observing a weaker correlation between the output gap and inflation, resulting in a flat estimated Phillips curve, would effect a quick disinflation because there is no short-run trade-off. This is exactly the opposite of the new Keynesian Phillips curve in Primiceri (2005): a flatter Phillips curve implies a more costly disinflation — observe the target criterion (47). For this reason Cogley and Sargent (2002) argue that during the 1970s a policymaker using a new classical Phillips curve should have detected the absence of trade-off fairly quickly. It becomes hard then to reconcile the late start of the Volcker disinflation with a new classical Phillips curve.

Cogley and Sargent (2005b) propose model uncertainty as a solution. The policymaker is concerned about robustness and chooses optimal policy using three alternative models for the evolution of inflation. The models are estimated recursively, and the posterior probabilities associated with each model computed every period. Optimal policy is determined by a Bayesian model-averaging procedure. The first model is of the form (46). The second includes inflation in first differences as the dependent variable, implying no long-run trade-off between inflation and output gap. Finally the third is a New Classical Phillips curve with no short-run trade-off, where unemployment is the dependent variable. This is also chosen to be the correct reference model.

In absence of a concern for robustness the central bank would quickly learn the correct model and disinflate quickly. However, disinflation would be very costly if the truth were in

fact represented by one of the first two models with a relatively flat Phillips curve. Indeed, Cogley and Sargent (2005b) show that the loss from full disinflation would be unbounded! Even a small weight on these alternative models would justify a cautious approach to disinflation. Only when all models generate finite costs of disinflation does the policymaker act aggressively toward inflation. These results highlight that even a policymaker committed to maximizing the right loss function can induce persistent sub-optimal outcomes as a result of imperfect knowledge about the correct structure of the economy.

Earlier results emphasize uncertainty and imperfect information about the private sector’s expectations formation process, while here private-sector expectations formation does not play as much of a role. Orphanides and Williams (2002, 2007, 2013), however, show the interaction between mismeasurement of natural rates and private sector learning adds a substantial degree of complexity to the central bank’s stabilization problem.⁴¹ Policymaker mistakes stemming from imperfect information about the economic environment are amplified by private-sector learning and forecasting problems.⁴² Policy conducted using optimal control methods, under the assumption of rational expectations and perfect information, can perform very poorly. Orphanides and Williams (2012) explain the rise of inflation during the 1970s with a combination of three main factors. First, the central bank assigned a relatively high weight on output gap stabilization, λ_x . Second, its estimate of the natural level of output y_t^n was too optimistic. Third, optimal policy was designed without taking into account the private sectors’ imperfect knowledge about the correct model of the economy. Optimal policy resulted in the great inflation.

A final account of the US inflation trends in the 1970s and 1980s, first pioneered by Sargent (1999), is explored in Cho, Williams, and Sargent (2002) and Sargent and Williams (2005). Adopting the same model described above, where the true data-generating process is a new classical Phillips curve, a parameter configuration is chosen to emphasize specific dynamics under learning. They show the economy converges to the neighborhood of the time-consistent equilibrium outcome with positive inflation, reflecting a perceived trade-off between inflation and unemployment. However, in response to certain sequences of shocks the economy can temporarily ‘escape’ to the optimal Ramsey, time-inconsistent, outcome. During these episodes the central bank learns the natural-rate hypothesis and proceeds to disinflate.⁴³

Result 7. *Imperfect information about natural rates and the structure of the model can lead to persistent deviations from policy objectives, despite full commitment by the policymaker.*

What kind of policy strategy could have avoided this outcome? Orphanides and Williams

⁴¹The analysis is based on a New Keynesian model with habit formation and price indexation where learning is implemented using the Euler Equation approach. The model is estimated on US data.

⁴²Although they do not focus on optimal policy, Lansing (2002) and Bullard and Eusepi (2005) emphasize that private-sector misperceptions about the natural rate and, more generally, about the model of the economy, were as important in explaining inflation in the 1970s as policymaker mistakes.

⁴³Details on the study of escape dynamics can be found in Williams (2014). A related theory proposed by Lubik and Matthes (2014) emphasizes indeterminacy of equilibrium induced by central bank learning, rather than escape dynamics.

(2007, 2013) show inertial policy rules, which respond more aggressively to inflation and less to economic activity than under rational expectations, are robust to both private-sector learning (characterized by different constant gains $\bar{g}$) and alternative degrees of mismeasurement in natural rates — which is related to the volatility of y_t^n and r_t^n , and the parameters in the updating rule (45). Policy rules with high inertia that respond to the growth rate of output gap (unemployment) by construction do not respond as much to variations in the estimated natural rate, avoiding the propagation of policy mistakes on expectations. Orphanides and Williams (2007) show inertial rules with such features can be justified as the outcome of a robust policy exercise using either Bayesian model averaging or a robust control approach. Moreover, Orphanides and Williams (2013) show with counterfactual simulations that, had the Federal Reserve followed such a rule, the great inflation would have been avoided and inflation expectations would have remained anchored through the 1970s and the 1980s. This result strengthens the claim that inertial rules are desirable under learning and imperfect information.

Summing up, the earlier results in section 6 rely on the central bank having correct knowledge of the agents' expectations formation mechanism. However, the discussion above shows these policy prescriptions are robust to some aspects of either policymaker uncertainty about the expectations formation process or the structural model of the economy.

Result 8. *Operational policy rules of the form (20) with $\rho \rightarrow 1$ are robust to imperfect knowledge about the model environment.*

7 Extensions

The Great Recession and concomitant policy response raise fundamental questions about the generality of the theory explicated above. Two omissions are notable: the role of fiscal policy in monetary policy design; and the consequences of the zero lower bound on nominal interest rates for the choice of monetary policy rule. This section redresses this, demonstrating both considerations represent important constraints on what monetary policy can achieve, but nonetheless, the basic lessons already learned are of greater importance in these environments. Before tackling these issues we discuss the role of central bank communication in improving the efficacy of monetary policy. While a topic of considerable interest in its own right, it serves to introduce some ideas relevant to subsequent analysis.

7.1 Central Bank Communication

A central bank's communication strategy is a crucial ingredient of inflation targeting. Inflation targeting central banks devote considerable effort to communicating information about monetary policy. This includes providing an assessment of the state of the economy, as well as numerous details on the projected evolution of different macroeconomic variables of interest, and how these projections inform the stance of policy. For example, many inflation targeters publish forecasts of inflation and real activity, where forecasts are conditioned on

‘appropriate policy’ which remains undisclosed. Notably, the Reserve Bank of New Zealand, the Riksbank and Norges Bank also present forecasts of their key policy rate.

While such practice offers basic accountability, it serves another purpose — to influence inflation expectations to achieve the most efficient short-run trade-off between output and inflation. Indeed, it is often argued that the greatest benefit from successful implementation of inflation targeting is the anchoring of expectations, with its stabilizing effect on macro-economic activity. Yet despite the importance of communication strategies, relatively little formal analysis in the context of dynamic stochastic general equilibrium models has been done on the mechanisms by which communication might prove beneficial.

An appealing feature of learning models is the property that subjective beliefs can differ from the objective probabilities implied by the economic model. This permits defining meaningfully the notion of anchored expectations as those beliefs consistent with the monetary policy strategy of the central bank. The possibility of beliefs being inconsistent with monetary policy strategy and, therefore, unanchored, presents a challenge for stabilization policy and permits examination of the role of communication in policy design. Communication can be modeled as providing households and firms with certain kinds of information about the monetary authority’s policy strategy. This information serves to simplify the forecasting problem of agents and to coordinate expectations in a desirable way.

To give content to these ideas recall the discussion of forecast-based instrument rules underpinning **Result 2**. Such rules encourage divergent learning dynamics: the central difficulty is agents, using a reduced-form model of inflation, income and nominal interest rates, project a path for real interest rates that is inconsistent with the monetary strategy of the central bank. Because forecasts for future real rates is the mechanism by which the central bank influences aggregate demand, stabilization policy is compromised. But what if the central bank conveys information about the adopted policy rule? Suppose the central bank announces policy will be implemented according to the simple rule (34) in all future periods. If this announcement is credible, knowledge of this rule simplifies agents’ forecasting problem. Rather than forecasting interest rates and inflation independently, forecasts can be made by combining reduced-form inflation forecasts with the policy rule to generate a policy-consistent forecast of the nominal interest rate. An advantage of this approach is forecasts of nominal interest rates will satisfy the Taylor Principle assuming $\phi_\pi > 1$. This confers advantages for stabilization policy. In fact, if agents make policy consistent forecasts then the Taylor principle is restored.

This basic insight is found in Eusepi (2005), Preston (2006) and developed further by Eusepi and Preston (2010). They propose various strategies of this kind, varying the types of information conveyed to private agents. In general, so long as the systematic component of policy is communicated, and understood to a reasonable degree of accuracy — for example agents have enough information to verify a central bank’s rule by use of statistical analysis — then this result continues to hold: communication improves the management of expectations.

Result 9. *If the central bank conveys enough information so that the private sector can infer its reaction function, then the Taylor Principle is restored.*

The result suggests publishing the forecast path for the policy rate is an important aspect

of a central bank communication strategy. The joint publication of inflation, economic activity and interest-rate forecasts, together with a fully articulated rationale behind the forecasts, allows agents to learn about the central bank reaction function. Indeed, the Norges Bank writes:

“As one of the reasons for publishing the forecast is to improve the general understanding of the Bank’s response pattern, it has been important to explain the logic behind the forecast. This includes shedding light on the considerations underlying each interest rate forecast and what their objectives and the trade-offs between them are”, Deputy Governor Jan F. Qvigstad, Limassol, 3 September 2010.

A range of papers explore related aspects of this basic idea. Orphanides and Williams (2005b) present a non microfounded model of inflation dynamics, where a central bank has direct control of inflation as the instrument of policy. While their model is always E-Stable, they show communicating the inflation target, so that agents need not estimate a drift in their belief structure, enhances inflation control and is welfare improving. Eusepi and Preston (2010) demonstrates that communicating the inflation target may not assist if shocks are persistent. Central then is an understanding of the systematic component of policy. Branch and Evans (2013) demonstrate in a global analysis of the New Keynesian model, under the Euler equation approach, conditions under which inflation beliefs approximate a random walk and are nearly self-fulfilling. Such beliefs hamper implementation of the central bank’s inflation target. Again, credible communication of the intended inflation target resolves instability. Similarly, Cogley, Matthes, and Sbordone (2015) consider the question of an optimal disinflation within a class of Taylor rules. They show agent uncertainty about the new policy can generate substantial instability leading to particular choices of monetary policy rule. Gibbs and Kulish (2015) develop a model of imperfect credibility of announcements. Less credible announcements, reflected in greater weight on historical patterns of data when forecasting, generate larger and more variable sacrifice ratios. Ascari, Florio, and Gobbi (2014), building on Eusepi (2005), show communication improves E-Stability in models where only one-period-ahead expectations matter, but where average inflation rates are greater than zero. Finally, Honkapohja and Mitra (2014) provide a global analysis of the zero lower bound. Communication greatly enhances stabilization properties of price-level rules. This is discussed in section 7.5.

Before moving on, a clear implication of the optimal policy exercises in section 6.3 is interest-rate beliefs impose a fundamental constraint on a central bank’s pursuit of price stability. Of course, in the simple example developed there, if agents understood that long-term interest rates were in fact constant in all periods, then the aggregate demand equation would no longer be a constraint. While highly simplistic, it does suggest communicating information about interest-rate policy, by providing projections of future interest-rate decisions, in addition to inflation and output, is highly desirable. Announced paths of inflation and output may be consistent with a range of beliefs about future interest rates — which may lead to instability.

To conclude, the efficacy of monetary policy depends on two related ingredients. First, the central bank’s commitment to the policy is credible. Second, private agents can properly

evaluate the effects of policy on economic decisions. While both these properties are intrinsic to a rational expectations equilibrium analysis, they need not be present in a learning analysis. When they are absent, shifting expectations can confound a central bank's effort to implement their objectives.

7.2 A Model with Debt

In contrast to a rational expectations analysis, Ricardian fiscal policy can compromise the efficacy of monetary policy under imperfect knowledge — even policies so far judged desirable. Consider a government seeking to finance a constant level of government purchases G by issuing government debt and levying lump-sum taxes, τ_t . Following Woodford (1998, 2001), assume the government issues one period debt, B_t^s , in zero supply, and long-dated debt, B_t^m , in positive supply with coupon payments $\rho^{T-(t+1)}$ for $T > t$ and $0 \leq \rho \leq 1$ and price P_t^m . The asset can be interpreted as a portfolio of infinitely many bonds, with weights along the maturity structure given by $\rho^{T-(t+1)}$ and average duration $(1 - \beta\rho)^{-1}$.

For simplicity of exposition, assume an endowment economy where individual income is determined by an exogenous i.i.d. process $y_t(i)$. This can be thought of as the flexible price limit of the model of section 2. Most results discussed below can be established with the endogenous supply of output and sticky prices. Eusepi and Preston (2013) demonstrate optimality, to a first-order approximation, requires

$$c_t(i) = (1 - \beta) y_t(i) - \sigma^{-1} \beta \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} (i_T - \pi_{T+1}) + \delta \left(b_{t-1}^m(i) - \pi_t + \rho \beta P_t^m + \beta \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} (i_T - \pi_{T+1} - (\beta^{-1} - 1) \tau_t(i)) \right) \quad (48)$$

where b_t^m denotes debt in real terms and all variables (with the usual slight abuse of notation) are expressed in log-deviations from their steady state values; δ is the steady-state structural surplus to GDP ratio, which is proportional to the debt-to-GDP ratio and σ^{-1} is the consumption intertemporal elasticity of substitution.⁴⁴ This consumption decision rule is a joint implication of first-order conditions governing the intertemporal allocation of consumption, the portfolio allocation decision across short- and long-term debt, and transversality.

Optimal decisions require forecasts of nominal interest rates, inflation and taxes into the indefinite future. The top line in (48) describes the evolution of consumption in absence of wealth effects from holding debt. This captures the standard transmission mechanism of monetary policy in the model under rational expectations, in the case of a Ricardian fiscal policy. The bottom line in (48), referred to as the ‘non-Ricardian’ component of consumption demand, measures the wealth effects from holding government debt net of taxes. It comprises three components. The first is the real value of debt holdings, the second measures the present value of real returns from holding debt (purchased in the current period) and the third component denotes the present value of taxes. In a rational

⁴⁴In general agents could have different holdings of public debt and different tax obligations. However, subsequent results assume these quantities are equal across households.

expectations equilibrium analysis, these terms sum precisely to zero — they constitute the intertemporal budget constraint of the government. Under imperfect knowledge, incorrect forecasts of returns and taxes imply the public debt is perceived as net wealth.⁴⁵ The model therefore has two competing and opposing channels through which monetary policy affects aggregate demand: the first is the usual intertemporal substitution of consumption (measured by the parameter σ^{-1}); the second is wealth effects from asset holding (measured by the parameter δ).

The remaining model equations specify the flow budget constraint of the government and the price of long debt, given by the expected present discounted value of all future one-period interest rates, where the discount factor is $\beta\rho$. The monetary and fiscal rules are

$$i_t = \phi_\pi \pi_t \text{ and } \tau_t = \phi_b b_t^m$$

and market clearing implies

$$\int_0^1 c_t(i) di = c_t = y_t \text{ and } \int_0^1 b_t^m(i) di = b_t^m.$$

The model is completed with a description of beliefs, which is deferred until after results under rational expectations are discussed.

7.3 The role of fiscal policy: Rational Expectations

The properties of the model under rational expectations are well understood. Leeper (1991) demonstrates two configurations of monetary and fiscal policy are consistent with a unique bounded rational expectations equilibrium in this class of model. Either monetary policy is active and fiscal policy is passive such that

$$\phi_\pi > 1 \text{ and } 1 < \phi_\tau < \frac{1 + \beta}{1 - \beta}$$

with inflation dynamics determined as $\pi_t = \phi_0 y_t$, for all t , giving the conventional Ricardian equilibrium; or monetary policy is passive and fiscal policy is active such that

$$0 \leq \phi_\pi < 1 \text{ and } 0 \leq \phi_\tau < 1 \text{ or } \phi_\tau > \frac{1 + \beta}{1 - \beta},$$

with inflation dynamics determined as $\pi_t = \phi_1 b_t^m + \phi_2 y_t$, giving a non-Ricardian equilibrium in which debt has monetary consequences — the fiscal theory of the price level. These results

⁴⁵Despite Ricardian Equivalence not being a general property of models with learning dynamics, the proposition does hold in several special cases of interest. When debt is in zero supply so that $\delta = 0$ then the non-Ricardian component of consumption demand is absent to a first-order approximation. Similarly, there will be no wealth effects on consumption demand if households hold certain beliefs that government debt will be retired by appropriate current and future tax policy. As noted by Evans, Honkapohja, and Mitra (2012), which examines the conditions for Ricardian Equivalence to hold in under both transition dynamics and convergence, beliefs must also satisfy certain conditions. Subsequent discussion will not impose these conditions. Further empirical work understanding fiscal beliefs would greatly assist evaluating the effects of fiscal policy in models of learning.

also hold in the benchmark New Keynesian model. Due to space constraints subsequent discussion under learning gives focus to the conventional assignment of policy. However, there is a growing body of evidence that the fiscal theory of the price level better characterizes certain episodes of US monetary history. Evans and Honkapohja (2007) and Eusepi and Preston (2011b) discuss stability of such equilibria under learning.

7.4 The role of fiscal policy: Learning Dynamics

For simplicity, assume agents know the policy regime in place, and that they need only forecast future average level taxes and inflation. The condition for E-Stability is

$$\phi_\pi > \max \left(1, \left[(2 - \beta) - \delta\sigma\beta \cdot \frac{(1 - \rho)\rho\beta}{1 - \rho\beta} \right]^{-1} \right).$$

Not only is the average level of indebtedness a constraint on monetary policy, but also the average duration of debt. As a result the Taylor principle does not hold. Consider two simple cases. First, if the wealth effects are relatively small, $\delta\sigma \approx 0$, then tax beliefs do not affect the dynamics of inflation. Dynamics under learning are then the same as in absence of debt: Eusepi and Preston (2013) show the E-Stability conditions are determined by the ODE (30) in section 4.3 and the Taylor principle applies. Second, if the government bonds have a duration of one period or they are all consol bonds (ρ takes the values of 0 or 1), then the evolution of taxes and debt again do not affect inflation beliefs. For these extreme values of average debt duration, the *quantity* of outstanding debt (and therefore taxes) does not respond to interest-rate expectations: the evolution of tax beliefs is not self-referential and the Taylor principle holds.⁴⁶ More generally, higher average levels of debt require stronger responses to inflation. In contrast, the average duration of debt induces a non-monotonic response from policy for stability — moderate durations of debt are most problematic.

While the discussion suggests it is safe to proceed without explicit treatment of fiscal variables when government debt is either of very short or very long maturity, this is not necessarily the case. Suppose monetary policy is given by the simple rule

$$i_t = \phi_\pi \hat{E}_{t-1} \pi_t$$

and that households *understand* this to be the case.⁴⁷ Analyze expectational stability in the neighborhood of the flexible-price equilibrium — formally $\alpha \rightarrow 0$. Then with one-period debt, the requirements for expectations stability are

$$1 < \phi_\tau < \frac{1 + \beta}{1 - \beta} \text{ and } \phi_\pi > \frac{1}{1 - \beta\delta\sigma}.$$

⁴⁶For more discussion on this point see Eusepi and Preston (2013). This result for $\rho = 0$ was first established by Evans and Honkapohja (2007) in a New Keynesian model with one-period ahead expectations and later extended by Eusepi and Preston (2012) to the anticipated utility framework, serving to extend **Result 1a** and **1b** to an explicit treatment of fiscal policy.

⁴⁷Absent knowledge of the policy rule **Result 2** continues to apply.

The Taylor principle is, again, no longer sufficient for stability. Fiscal policy again constrains the class of monetary policy rule consistent with stability. The higher the average level of indebtedness, the more aggressive must be monetary policy to stabilize inflation. The hypothetical extreme $\delta = \sigma = 1$ recovers the necessary condition for stability identified by Preston (2006) and discussed in section 5.1. In this case, however, the policy rule is fully understood by the public! This highlights the self-referentiality of interest-rate policy and interest-rate beliefs is a fundamental source of instability and more important when fiscal policy is introduced.⁴⁸

Result 10. *In economies with high government debt, wealth effects from government liabilities can reduce the effectiveness of monetary policy in controlling inflation. Furthermore, the average maturity of debt also influences monetary policy outcomes.*

Fiscal policy can frustrate the pursuit of price stability. Eusepi and Preston (2013) develop this insight to quantify the importance of sound fiscal policy for macroeconomic stability in the US. They show using counterfactual simulations that had fiscal policy been less responsible, reflected in levels, and average duration, of debt of the kind commonly observed after the financial crisis, the great moderation would not have been observed — the average volatility of inflation and output would have been comparable to the 1960s and 1970s.⁴⁹ This signals potential policy challenges ahead for many advanced economies.

7.5 Zero Lower Bound and global dynamics

The results discussed in the previous sections ignore a key constraint on monetary policy: the zero lower bound (ZLB) on the short-term nominal rate. The ZLB severely limits the power of monetary policy even under perfect information. While it is impossible to do justice to this substantial literature, there are two basic concerns for the discussion here. First, Eggertsson and Woodford (2003) show large and persistent shocks that drive the economy to the ZLB for a protracted period of time can generate large output losses driven by a deflationary spiral. Second, Benhabib, Schmitt-Grohe, and Uribe (2001) demonstrate commitment to policy rules of the form (20), generally delivers local determinacy but global indeterminacy: among possible equilibria, there exist a continuum of equilibrium paths converging to a deflationary steady state. These observations represent an important challenge to monetary policy design.

Putting aside the matter of global indeterminacy, policy rules that respond to the price level, and specifically induce a response to the lagged price level, have desirable properties under rational expectations. Eggertsson and Woodford (2003) show a price-level targeting

⁴⁸Interestingly, in absence of central bank communication of the policy rule, the stance of fiscal policy is irrelevant. The instability stemming from policy expectations being disconnected from the actual policy rule dominates.

⁴⁹In related analysis, Sinha (2015) demonstrates that identified debt shocks from simulated data from a closely related model explain 30 percent of the rise in short-term yields and 10 percent of the rise in long-term yields identified using a vector auto-regression.

rule (14) comes very close to implementing the optimal policy taking into account the ZLB.⁵⁰ Moreover, they display robustness to model uncertainty under learning — recall section 5.3. How do they perform when subject to the constraint of the ZLB? Williams (2010) considers a simple reduced-form New Keynesian model with learning where agents have to form multi-period forecasts of the output gap, inflation and the nominal interest rate. Using simulation exercises, the study evaluates the performance of policy rules incorporating a response to the price-level of the form

$$i_t = \max \left[r_t^n + \phi_p (\ln p_t - \ln p_t^*) + \phi_{\Delta x} (x_t - x_{t-1}) + \phi_x x_t, 0 \right] \quad (49)$$

where $\ln p_t^* = \ln p_{t-1}^* + \pi^*$. Agents' knowledge of the policy rule is key to its effectiveness. In absence of central bank communication, or, equivalently, in absence of central bank credibility, the economy performs quite poorly under learning. The economy is frequently at the ZLB (five times more frequently than under rational expectations), which comes at the cost of sharp reductions in output. The frequency of these episodes is reduced substantially if the public understands the policy rule and forms policy-consistent forecasts of nominal interest rates. In addition, the effectiveness of policy can be further enhanced by allowing the policy rule to be more aggressive in response to deviations of the price level from target than is required under rational expectations. Each of these properties are desirable in absence of the ZLB constraint.

Of course, policy rules of this kind are not immune from global indeterminacy under rational expectations. Focusing purely on the local stability of rational expectations equilibria, it would be tempting to dismiss the practical relevance of a liquidity trap steady state (and the paths converging to it) as it is generally an equilibrium unstable under learning — see Evans, Guse, and Honkapohja (2008) and Eusepi (2007). However, a careful analysis of the global dynamics of the economy under learning reveals that, despite being unstable under learning, the liquidity trap equilibrium has a strong influence on the global behavior of the economic system. To anticipate some results, large shocks can push the economy into a prolonged period of low or negative inflation and low levels of output, very close to the 'liquidity trap' observed in Japan in the past decade. The perils of deflationary traps have been an on-going policy concern in the aftermath of the global financial crisis.⁵¹

A number of studies analyze global dynamics under Taylor-type policy rules in the simple New Keynesian model, under both finite-horizon and infinite-horizon decision rules.⁵² Regardless of which nonlinear decision rules are employed, these papers focus on a deterministic set up where agents believe the equilibrium is a possibly-drifting noisy steady state, the object about which they learn. The economic system is described by the following true data-generating process for output, inflation and the nominal interest rate

$$z_t = \mathbb{H}(\omega_{t-1})$$

⁵⁰However, Reifschneider and Williams (2000) also adduce evidence that a Taylor-type rule with a high degree of interest smoothing and a policy of price-level targeting are not equivalent in the presence of ZLB. Such rules prescribe excessively tight policy because of their dependence of past interest rates.

⁵¹See, for example, Bullard (2010).

⁵²Evans and Honkapohja (2005), Evans, Guse, and Honkapohja (2008), Evans and Honkapohja (2009), Eusepi (2010) and Honkapohja and Mitra (2014). These models assume Rotemberg pricing, which has the same first-order implications as Calvo pricing but it is more convenient for nonlinear analysis.

and the constant-gain rule

$$\omega_t = \omega_{t-1} + \bar{g} (\mathbb{G}(\mathbb{H}(\omega_{t-1})) - \omega_{t-1})$$

where ω_t is a vector including the estimated steady-state values and $\mathbb{H}(\cdot), \mathbb{G}(\mathbb{H}(\cdot))$ are nonlinear functions — details on the analysis of such systems can be found in Evans and Honkapohja 2001. Following a different approach, Bullard and Cho (2005) consider a linear economy described by (31) and (32) where private agents learn with a constant-gain algorithm and the policymaker follows a simple Taylor rule satisfying the Taylor principle. Their economy fluctuates most of the time around the locally stable inflation target steady-state equilibrium. However, certain sequences of shocks can trigger the economy to ‘escape’ towards a period of persistently low interest rates and deflation, observationally equivalent to a liquidity trap, before reverting back to fluctuating around the equilibrium.⁵³

These studies show the existence of a ‘corridor of stability’: small shocks to expectations deliver convergence back to the inflation target.⁵⁴ In contrast, sufficiently large shocks induce a deflationary spiral preventing convergence back to equilibrium. And even within the corridor of stability, sufficiently large shocks can lead to a prolonged period of low inflation and output before convergence occurs. **Figure 9** illustrates possible paths for output and inflation in a simple New Keynesian model under a Taylor rule.⁵⁵ The figure replicates the phase diagram generated by the model of Eusepi (2011). Agents make decisions according to the Euler equation approach. However, because of information delays, in absence of central bank communication they have to forecast one-step-ahead the current interest rate, as well as output and inflation. Output is expressed in percentage deviations from its steady-state value (at the inflation target steady state), while inflation is expressed in annualized percentage terms. Models without informational delays, or in which decision rules imply forecasts over an infinite horizon, show similar qualitative behavior of output and inflation — see, for example, Evans, Guse, and Honkapohja (2008), Evans and Honkapohja (2010) and Benhabib, Evans, and Honkapohja (2014) which are discussed below.

The steady state corresponding to the inflation target is a sink (implying local stability under learning) while the liquidity trap steady state is a saddle (unstable under learning). Interestingly the stable (dashed line) and unstable (solid line) manifolds associated to the liquidity trap steady-state ‘wrap’ the inflation target steady state, defining the corridor of stability in the economy. This corridor is better understood by considering three different initial conditions for inflation and output. If the economy suffers a small shock (point A), then, by the local properties of the Taylor principle, the policy rule induces convergence back to the steady state. A shock that is sufficiently large, but within the boundaries of the dashed line (point B), leads to a period of deflation, as the economy is attracted by the liquidity trap equilibrium. As the economy gets further away from the inflation target equilibrium,

⁵³The mechanism is the same as described in Sargent (1999) and section 6.4. One key assumption to obtain escapes in the model is that the inflation target set by the central bank is in fact the inflation target estimated by the private sector in real time, so that the nominal anchor of the system can slowly drift in response to shocks.

⁵⁴The term ‘corridor of stability’ dates back to Leijonhufvud (1973), which discusses the stabilizing effects of market forces in response to shocks.

⁵⁵Again, the Taylor rule satisfies the Taylor principle at the inflation target equilibrium.

monetary policy becomes less effective because the policy rule violates the Taylor principle near the liquidity trap equilibrium. However, the liquidity trap equilibrium is unstable and the economy slowly reverts back to its steady state. What are the forces behind this recovery? After a prolonged period of deflation and slow output growth, the low interest rate ultimately stimulates spending and the economy reverts back towards the inflation target equilibrium.⁵⁶ Finally, a sufficiently large shock (point C) triggers a deflationary spiral with no return to equilibrium. In this case, the deflationary impetus in expectations is too strong to be offset by stimulative monetary policy, and the economy would eventually violate feasibility constraints. Along these deflationary paths either the private sector or policymakers would change their behavior. The most likely scenario involves a change in policy. In fact, when considering such large shocks, both assumptions of a time-invariant policy rule and a time-invariant learning algorithm lack realism.

Result 11a. *Under Taylor rules which satisfy locally the Taylor principle, large shocks can drive expectations and the economy to a prolonged period of deflation and low output.*

As discussed in section 7.1, agents' knowledge of the policy rule has important implications for policy effectiveness. **Figure 10** shows the phase diagram of the model where agents understand the policy rule and therefore only have to forecast inflation and output gap. The stability corridor enlarges considerably under a regime of policy communication. All initial conditions lead to convergence to the inflation target. Honkapohja and Mitra (2014) discuss a similar result in the case of a policy rule that responds to the price level (or nominal output), such as (49). They show that without any guidance about the policy rule, the corridor of stability under a policy rule responding to the price level is considerably narrower than a rule responding to the inflation rate. Furthermore, the inflation target steady state is less robust to higher constant gains under price-level rules than under inflation-based policy rules. Conversely, with some guidance about the policy regime — specifically the desired price path — the stability properties of price-level rules improve substantially, even relative to a conventional inflation-based rule.⁵⁷

Result 11b. *Lack of communication increases the likelihood of entering a deflationary path. The stability corridor is smaller under a policy rule responding to the price level.*

Effective communication of policy does not eliminate the perils of Taylor rules. Although weaker, the influence of the liquidity trap equilibrium persists. For example, initial conditions in C continue to engender a period of persistent disinflation and low output. Section 7.2 discussed the interactions of monetary policy and government debt under learning. In the ZLB results discussed so far, government liabilities play no direct role because either: i) household and government intertemporal budget constraints are not taken into account; or

⁵⁶Notice that the economy overshoots its steady state before convergence.

⁵⁷However, even with some guidance about the policy rule, (49) is shown to induce local instability for relatively small values of the constant gain.

ii) agents in the economy fully understand that fiscal policy is Ricardian. Benhabib, Evans, and Honkapohja (2014) investigate whether wealth effects from government liabilities can prevent a deflationary outcome via the Pigou effect: falling prices increase perceived wealth and aggregate demand, boosting the economy out of the liquidity trap. Numerical examples are given where the wealth effects are strong enough to ensure convergence. An initial state of the economy that would lead to a deflationary path in absence of wealth effects on aggregate demand (for example, point C in **figure 9**) is in fact consistent with convergence to the inflation target steady state once wealth effects on aggregate demand are considered.⁵⁸

Results 11c. *Under a passive fiscal regime, perceived wealth effects from government liabilities can significantly enlarge the stability corridor.*

Pigou effects, however, produce convergence involving overshooting of output and inflation, so that alternative fiscal and monetary measures might be required to further stabilize the economy. Evans and Honkapohja (2010) evaluate the effectiveness of a temporary switch from a standard Taylor rule to a near-zero constant interest rate when inflation crosses a threshold (above that implied by the liquidity trap steady state). Independently of the duration of the policy, the stability corridor of the economy is not significantly altered, even if the policy is fully understood by the public. However, coupling the temporary interest-rate peg with a commitment to increase government spending to whatever level needed to keep inflation from falling below the threshold can eliminate the influence of the liquidity trap equilibrium. This kind of policy, however, delivers a fairly volatile economic environment in the transition to the inflation target steady-state. The analysis also begs the question of the size of fiscal expansion required to keep inflation from falling, in response to a negative shock similar to the recent financial crisis.

Result 11d. *A combination of expansionary fiscal and monetary policies that are well understood by the public can, if properly designed, eliminate deflationary equilibria.*

As for many topics discussed in this survey, a proper evaluation of this claim requires considerably more quantitative and robustness analysis. Regardless, the results convey the general idea that when constrained by the ZLB, policy can steer expectations even under imperfect information about the economic environment, where agents update their view of the economy according to simple time-invariant updating algorithms. As in environments where the ZLB is not a constraint, appropriately designed inertial policy rules that are clearly communicated have considerable merit.

8 Further directions

At a broad level much more needs to be understood about the classes of beliefs that determine actual firm and household behavior. Further study and development of survey data, financial market data and experimental data would be immensely helpful in this regard.

⁵⁸The result depends on fiscal policy being passive in the sense of Leeper (1991).

With this knowledge in hand, further quantitative and empirical research would greatly assist policy design in environments where agents might not fully comprehend the specific details of the policy regime in place. More specifically, there are at least four dimensions where the literature could grow. Two dimensions are specific to the learning literature. Two others are common to the New Keynesian literature in general.

First, expectations formation depends on the choice of the gain sequence, g_t . In this survey, as in most of the underlying literature it reflects, we have taken the gain as an exogenous parameter. However, this variable might well be an important choice variable for agents. The chosen gain could be a function of, among other factors, the policy regime in place and therefore subject to the Lucas critique — see, for example, Evans and Ramey (2006). As such, it should be integral part of the policy optimization problem. A few papers in the literature treat the constant gain as a decision variable. In Evans and Honkapohja (1993) and Marcet and Nicolini (2003) the equilibrium gain is the gain that minimizes the forecast mean-square error of each individual agents: that is, no agent has an incentive to choose a different gain in equilibrium. Marcet and Nicolini (2003) discuss the emergence of hyperinflationary episodes in response to seigniorage needs, and show that a change in policy leads to a change in the equilibrium constant gain. It is therefore critical to evaluate the policies discussed above by letting the constant gain be an endogenous variable in the model. Moreover, assuming either a constant gain or a decreasing gain is unnecessarily restrictive. A plausible learning algorithm could entail an endogenous switch between different gains, depending perhaps on the forecasting performance of the implied beliefs. Building on ideas first proposed by Brock and Hommes (1997), Marcet and Nicolini (2003), Evans, Honkapohja, Sargent, and Williams (2012) and Milani (2014) implement versions of such an algorithm.⁵⁹

The second dimension where the literature could and should expand is to abandon the assumption of homogeneous expectations. As documented by Branch (2004), Mankiw, Reis, and Wolfers (2004) and Andrade, Crump, Eusepi, and Moench (2013), professional forecasters disagree both in their short- and long-term forecasts of inflation, economic activity and the policy rate. In addition central bank forecasts are often not in agreement with the outlook of market participants. The efficacy of monetary policy might be compromised by heterogeneous expectations and will undoubtedly introduce distributional considerations. A few papers have investigated the consequences of heterogeneous private beliefs for monetary policy at a theoretical level — Branch (2004), Branch and McGough (2009, 2010, 2011). Honkapohja and Mitra (2005) and Preston (2008) explore disagreement between private sector and central bank forecasts. In an interesting study, Assenza, Heemeijer, Hommes, and Massaro (2014) develop laboratory experiments designed to evaluate performance of Taylor-type rules in presence of expectational heterogeneity. More should be done to develop testable theoretical predictions that can be confronted with survey and other data on expectations.

Third, further quantitative assessment of learning for macroeconomic dynamics and policy design would be invaluable. There is a growing evidence that learning models can explain features of data that are deemed a puzzle when viewed through the lens of a rational

⁵⁹See also Gaus (2014) and Kostyshyna (2012) for alternative time-varying gain algorithms.

expectations analysis. While most central banks use state of the art DSGE models of the kind proposed by Smets and Wouters (2007) for policy analysis, the large majority retains the assumption of rational expectations. More should be done to conduct policy evaluation in medium-scale New Keynesian models with learning. At the forefront of this effort are Milani (2004, 2008, 2011, 2012) and Sloboydan and Wouters (2012a, 2012b) which use Bayesian methods to estimate small- and medium-scale DSGE models. With the exception of Eusepi and Preston (2013), these studies adopt the Euler equation approach, so that some of the aspects of the optimal policy problem discussed in this survey cannot be studied. Moreover, as in the case of forecast disagreement, more use should be made of survey-based forecasts to discipline the class of permitted learning algorithms employed in estimation — for some use of survey-data see, for example, Milani (2011), Eusepi and Preston (2011, 2013), Sloboydan and Wouters (2012b), Ormeno and Molnar (2013).

Lastly, the recent financial crisis has stimulated efforts to introduce financial frictions and asset price dynamics in the baseline New Keynesian framework. It is easy to argue that such innovations could have important implications for the design of optimal monetary policy. These considerations are more compelling for models that incorporate learning, where asset prices, financial frictions and beliefs dynamics can deliver substantial variations from the rational expectations paradigm. Indeed, there has been a growing literature that attempts with a certain degree of success to explain asset prices with learning dynamics.⁶⁰ One more challenge for the literature is to integrate these models into the standard New Keynesian model to study their implication for monetary policy.⁶¹

9 Conclusions

The main principles of rational expectations analysis of monetary policy survive the introduction of imperfect information and learning but are modified, in some cases substantially. Not surprisingly, the science of monetary policy becomes more complex under imperfect knowledge, as learning can impose substantial limits to the effectiveness of monetary policy. Importantly central bank communication is shown to be one of the most important determinants of policy success, along with policy procedures that induce an appropriate dependence of current interest rates on private sector expectations, and an appropriate form of history dependence. We suggest re-writing the principles described in section 2 as follows:

1. Optimal policy under learning is complicated by multiple intertemporal trade-offs that are absent under rational expectations. As a result, the Divine Coincidence does not generally hold.

⁶⁰Contributions include Adam, Beutel and Marcet (2013), Adam, Marcet, and Nicolini (2012), Carceles-Poveda and Giannitsarou (2008), Benhabib and Dave (2014), Branch and Evans (2010, 2011), and Sinha (2014).

⁶¹Recent studies on the interactions between monetary policy and asset pricing under learning include Airaudo, Nisticò, and Zanna (2014), Winkler (2015), Gelain, Lansing, and Mendicino (2013), Kitney (2014), Pintus and Suda (2014), Gelain, Lansing, and Natvik (2015), Dewachter, Iania, and Lyrio (2011), Milani (2008) and Caputo, Medina, and Soto (2011).

2. In the neighborhood of rational expectations, optimal policy under learning is generally more aggressive towards inflation than under rational expectations. This is due both to the intertemporal trade-offs induced by the expectations formation mechanism and because of non-Ricardian wealth effects generated by the dynamics of government debt. When beliefs are sufficiently far from rational expectations equilibrium, optimal policy stipulates less aggressive responses to changing economic conditions than under rational expectations.
3. Optimal policy under learning induces history dependence which improves the short-run stabilization trade-off between inflation and the output gap. Such history dependence is however different in nature to the optimal commitment solution and depends on agents' learning rules.
4. Policy rules that either respond to the price level, or incorporate interest-rate smoothing, have good stabilization properties. When subject to the zero lower bound constraint the efficacy of policy depends on effective and credible communication.
5. Simple policy rules that incorporate a response to the price-level or a sufficient degree of interest-rate smoothing exhibit a high degree of robustness to model uncertainty — reflecting mis-measured variables (natural rates) or different assumptions about the expectations formation mechanism of private agents.

The body of evidence supports the adoption of some form of price-level targeting. Provided such an approach is credible, well communicated and understood by the public, policy rules consistent with PLT provide the desired aggressive stance toward inflation, and also appropriate history dependence. They also insure a better performance once the ZLB is taken into account by reducing the probability that the economy hits the ZLB, and better management of expectations when this constraint is binding. Importantly, failure to properly communicate these rules can imply poor performance.⁶² An implication is that switching to some form of price-level targeting is most desirable under economic circumstances when the policymaker has the tools to establish credibility.⁶³ For example, switching to a PLT regime while a the zero lower bound might not have the desired impact on expectations and therefore might not deliver a faster exit to positive interest rates.

References

ADAM, K., J. BEUTEL, AND A. MARCET (2013): “Stock Price Booms and Expected Capital Gains,” unpublished, University of Mannheim.

⁶²This resonates with Woodford (2010), which shows that price-level targeting is robust to uncertainty about expectation formation. Here, the deviation from rational expectation is modelled following the robust control literature.

⁶³On the costs of introducing price level targeting under imperfect credibility see Kryvtsov, Shukayev, and Ueberfeldt (2008) and Cateau, Kryvtsov, Shukayev, and Ueberfeldt (2009).

- ADAM, K., G. W. EVANS, AND S. HONKAPOHJA (2006): “Are hyperinflation paths learnable?,” *Journal of Economic Dynamics and Control*, 30(12), 2725–2748.
- 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.-P. NICOLINI (2012): “Sock Market Volatility and Learning,” unpublished, University of Mannheim.
- ADAM, K., AND M. PADULA (2011): “Inflation Dynamics and Subjective Expectations in the United States,” *Economic Inquiry*, 49(1), 13–25.
- AIRAUDO, M., S. NISTICÒ, AND L.-F. ZANNA (2014): “Learning, Monetary Policy and Asset Prices,” Working Papers 4/14, Sapienza University of Rome, DISS.
- ANDRADE, P., R. K. CRUMP, S. EUSEPI, AND E. MOENCH (2013): “Noisy Information and Fundamental Disagreement,” Federal Reserve Bank of New York Staff Reports, Numer 655.
- ASCARI, G., A. FLORIO, AND A. GOBBI (2014): “Transparency, Expectations Anchoring and the Inflation Target,” unpublished, University of Oxford.
- ASSENZA, T., P. HEEMEIJER, C. H. HOMMES, AND D. MASSARO (2014): “Managing Self-organization of Expectations through Monetary Policy: a Macro Experiment,” unpublished, University of Amsterdam.
- BATINI, N., AND A. G. HALDANE (1999): “Forward-Looking Rules for Monetary Policy,” in *Monetary Policy Rules*, ed. by J. Taylor. University of Chicago Press, Chicago.
- BENHABIB, J., AND C. DAVE (2014): “Learning, Large Deviations and Rare Events,” *Review of Economic Dynamics*, 17(3), 367–382.
- BENHABIB, J., G. W. EVANS, AND S. HONKAPOHJA (2014): “Liquidity traps and expectation dynamics: Fiscal stimulus or fiscal austerity?,” *Journal of Economic Dynamics and Control*, 45(C), 220–238.
- BENHABIB, J., S. SCHMITT-GROHE, AND M. URIBE (2001): “The Perils of Taylor Rules,” *Journal of Economic Theory*, 96(1-2), 40–69.
- BERARDI, M., AND J. K. GALIMBERTI (2013): “A note on exact correspondences between adaptive learning algorithms and the Kalman filter,” *Economics Letters*, 118, 139–142.
- BERNANKE, B. S., AND M. WOODFORD (1997): “Inflation Forecasts and Monetary Policy,” *Journal of Money, Credit, and Banking*, 29(4), 653–686.
- BLANCHARD, O., AND J. GALÍ (2007): “Real Wage Rigidities and the New Keynesian Model,” *Journal of Money, Credit and Banking*, 39(s1), 35–65.
- BOMFIM, A., R. J. TETLOW, P. VON ZUR MUEHLEN, AND J. WILLIAMS (1997): “Expectations, Learning and the Costs of Disinflation: Experiments using the FRB/US Model,” manuscript, Federal Reserve Board of Governors.
- BRANCH, W., AND B. MCGOUGH (2011): “Business cycle amplification with heterogeneous expectations,” *Economic Theory*, 47(2), 395–421.

- BRANCH, W. A. (2004): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *Economic Journal*, 114(497), 592–621.
- BRANCH, W. A., T. DAVIG, AND B. MCGOUGH (2013): “Adaptive Learning In Regime-Switching Models,” *Macroeconomic Dynamics*, 17(05), 998–1022.
- BRANCH, W. A., AND G. W. EVANS (2006): “A Simple Recursive Forecasting Model,” *Economics Letters*, (91), 158–166.
- (2010): “Asset Return Dynamics and Learning,” *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.
- (2013): “Unstable Inflation Targets,” manuscript, University of California, Irvine.
- BRANCH, W. A., AND B. MCGOUGH (2009): “A New Keynesian model with heterogeneous expectations,” *Journal of Economic Dynamics and Control*, 33(5), 1036–1051.
- (2010): “Dynamic predictor selection in a new Keynesian model with heterogeneous expectations,” *Journal of Economic Dynamics and Control*, 34(8), 1492–1508.
- BRAY, M. (1982): “Learning, Estimation and the Stability of Rational Expectations,” *Journal of Economic Theory*, 26, 318–339.
- BRAY, M., AND N. E. SAVIN (1986): “Rational Expectations Equilibria, 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), 1059–1096.
- BULLARD, J., AND I.-K. CHO (2005): “Escapist policy rules,” *Journal of Economic Dynamics and Control*, 29(11), 1841–1865.
- BULLARD, J., AND S. EUSEPI (2005): “Did the Great Inflation Occur Despite Policymaker Commitment to a Taylor Rule?,” *Review of Economic Dynamics*, 8(2), 324–359.
- (2014): “When Does Determinacy Imply Expectational Stability?,” *International Economic Review*, 55, 1–22.
- 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., AND E. SCHALING (2005): “Monetary Policy, Determinacy and Learnability in the Open Economy,” Discussion paper.
- BULLARD, J., AND A. SINGH (2008): “Worldwide macroeconomic stability and monetary policy rules,” *Journal of Monetary Economics*, 55(Supplemen), S34–S47.
- BULLARD, J. B. (1992): “Time-varying parameters and nonconvergence to rational expectations under least square learning,” *Economics Letters*, 40, 159–166.

- (2006): “The learnability criterion and monetary policy,” *Review*, (May), 203–217.
- CAPUTO, R., J. P. MEDINA, AND C. SOTO (2011): “The Financial Accelerator under Learning and the Role of Monetary Policy,” in *Monetary Policy under Financial Turbulence*, ed. by L. F. Céspedes, R. Chang, and D. Saravia, vol. 16 of *Central Banking, Analysis, and Economic Policies Book Series*, chap. 7, pp. 185–218. Central Bank of Chile.
- CARCELES-POVEDA, E., AND C. GIANNITSAROU (2008): “Asset Pricing with Adaptive Learning,” *Review of Economic Dynamics*, 11(3), 629–651.
- CATEAU, G., O. KRYVTSOV, M. SHUKAYEV, AND A. UEBERFELDT (2009): “Adopting Price-Level Targeting Under Imperfect Credibility in ToTEM,” Working Paper 2009-17, Bank of Canada.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): “Escaping Nash Inflation,” *Review of Economic Studies*, 69(1), 1–40.
- CHO, S. (2014): “Determinacy and E-Stability under Contemporaneous Expectations,” unpublished, Yonsei University.
- CLARIDA, R., J. GALI, AND M. GERTLER (1998): “Monetary Policy Rules in Practice: Some International Evidence,” *European Economic Review*, 42, 1033–1067.
- (1999): “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37, 1661–1707.
- (2000): “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, 115, 147–180.
- COGLEY, T., C. MATTHES, AND A. M. SBORDONE (2015): “Optimized Taylor rules for disinflation when agents are learning,” *Journal of Monetary Economics*, 72(C), 131–147.
- COGLEY, T., G. E. PRIMICERI, AND T. J. SARGENT (2010): “Inflation-Gap Persistence in the US,” *American Economic Journal: Macroeconomics*, 2(1), 43–69.
- COGLEY, T., AND T. J. SARGENT (2002): “Evolving Post-World War II U.S. Inflation Dynamics,” in *NBER Macroeconomics Annual 2001*, pp. 331–379.
- (2005a): “The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty,” *Review of Economic Dynamics*, 8(2), 528–563.
- (2005b): “The conquest of US inflation: Learning and robustness to model uncertainty,” *Review of Economic Dynamics*, 8(2), 528–563.
- (2008): “Anticipated Utility And Rational Expectations As Approximations Of Bayesian Decision Making,” *International Economic Review*, 49(1), 185–221.
- COGLEY, T., AND A. M. SBORDONE (2008): “Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve,” *American Economic Review*, 98(5), 2101–26.
- CRUMP, R., S. EUSEPI, AND E. MOENCH (2015): “The Term Structure of Expectations and Bond Yields,” unpublished, Federal Reserve Bank of New York.

- DENNIS, R., AND F. RAVENNA (2008): “Learning and Optimal Monetary Policy,” *Journal of Economic Dynamics and Control*, 32, 1964–1994.
- DEWACHTER, H., L. IANIA, AND M. LYRIO (2011): “A New-Keynesian model of the yield curve with learning dynamics: A Bayesian evaluation,” MPRA Paper 34461, University Library of Munich, Germany.
- DUFFY, J., AND W. XIAO (2011): “Investment and Monetary Policy: Learning and Determinacy of Equilibrium,” *Journal of Money, Credit and Banking*, 43(5), 959–992.
- EDGE, R., T. LAUBACH, AND J. WILLIAMS (2007): “Learning and Shifts in Long-Run Productivity Growth,” *Journal of Monetary Economics*, pp. 2421–2438.
- EGGERTSSON, G., AND M. WOODFORD (2003): “The Zero Bound on Interest Rates and Optimal Monetary Policy,” .
- ELLISON, M., AND J. PEARLMAN (2011): “Saddlepath learning,” *Journal of Economic Theory*, 146(4), 1500–1519.
- EUSEPI, S. (2005): “Central Bank Transparency under Model Uncertainty,” FRBNY Working Paper.
- (2007): “Learnability and Monetary Policy: A Global Perspective,” *Journal of Monetary Economics*, (54), 1115–1131.
- (2010): “Central Bank Communication and the Liquidity Trap,” *Journal of Money, Credit and Banking*, 42(2-3), 373–397.
- EUSEPI, S., M. GIANNONI, AND B. PRESTON (2012): “Long-Term Debt Pricing and Monetary Policy Transmission Under Imperfect Knowledge,” Federal Reserve Bank of New York Staff Reports, Number 547.
- (2015): “The Limits of Monetary Policy with Long-term Drift in Expectations,” unpublished, The University of Melbourne.
- EUSEPI, S., AND B. PRESTON (2010): “Central Bank Communication and Macroeconomic Stabilization,” *American Economic Journal: Macroeconomics*, 2, 235–271.
- (2011a): “Expectations, Learning and Business Cycle Fluctuations,” *American Economic Review*, 101(6), 2844–2872.
- (2011b): “Learning the fiscal theory of the price level: Some consequences of debt-management policy,” *Journal of the Japanese and International Economies*, 25(4), 358–379.
- (2012): “Debt, Policy Uncertainty and Expectations Stabilization,” *Journal of the European Economics Association*, 10(4), 860–886.
- (2013): “Fiscal Foundations of Inflation: Imperfect Knowledge,” unpublished, The University of Melbourne.
- EVANS, G., AND B. MCGOUGH (2014): “Learning to Optimize,” manuscript, University of Oregon.

- EVANS, G. W., E. GUSE, AND S. HONKAPOHJA (2008): “Liquidity traps, learning and stagnation,” *European Economic Review*, 52(8), 1438–1463.
- EVANS, G. W., AND S. HONKAPOHJA (1994): “Learning, Convergence, and Stability with Multiple Rational Expectations Equilibria,” *European Economic Review*, 38, 1071–1098.
- (1998): “Economic Dynamics with Learning: New Stability Results,” *The Review of Economic Studies*, 1, 23–44.
- (2001): *Learning and Expectations in Economics*. Princeton, Princeton University Press.
- (2003a): “Adaptive learning and monetary policy design,” *Proceedings*, pp. 1045–1084.
- (2003b): “Expectations and the Stability Problem for Optimal Monetary Policies,” *Review of Economic Studies*, 70(4), 807–824.
- (2005): “Policy Interaction, Expectations and the Liquidity Trap,” *Review of Economic Dynamics*, 8, 303–323.
- (2006): “Monetary Policy, Expectations and Commitment,” *Scandinavian Journal of Economics*, 108, 15–38.
- (2007): “Policy Interaction, Learning and the Fiscal Theory of Prices,” *Macroeconomic Dynamics*, 11, 665–690.
- (2009a): “Expectations, Learning and Monetary Policy: An Overview of Recent Research,” in *Monetary Policy under Uncertainty and Learning*, ed. by K. Schmidt-Hebbel, C. E. Walsh, N. L. S. Editor), and K. S.-H. (Series, vol. 13 of *Central Banking, Analysis, and Economic Policies Book Series*, chap. 2, pp. 027–076. Central Bank of Chile.
- (2009b): “Learning and Macroeconomics,” *Annual Review of Economics*, 1(1), 421–451.
- (2009c): “Robust Learning Stability with Operational Monetary Policy Rules,” in *Monetary Policy under Uncertainty and Learning*, ed. by K. Schmidt-Hebbel, and C. Walsh, pp. 27–76.
- (2010): “Expectations, Deflation Traps and Macroeconomic Policy,” in *Twenty Years of Inflation Targeting: Lessons Learned and Future Prospects*, ed. by D. Cobham, O. Eitheim, S. Gerlach, and J. F. Qvigstad, pp. 232–260.
- 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.
- (2013): “Notes on Agents’ Behavioral Rules under Adaptive Learning and Studies of Monetary Policy,” in *Macroeconomics at the Service of Public Policy*, ed. by T. J. Sargent, and J. Vilmunen. Oxford University Press.

- EVANS, G. W., S. HONKAPOHJA, T. SARGENT, AND N. WILLIAMS (2012): “Bayesian Model Averaging, Learning and Model Selection,” CDMA Working Paper Series 201203, Centre for Dynamic Macroeconomic Analysis.
- EVANS, G. W., AND B. MCGOUGH (2007): “Optimal Constrained Interest-Rate Rules,” *Journal of Money, Credit and Banking*, 39(6), 1335–1356.
- EVANS, G. W., AND G. RAMEY (2006): “Adaptive expectations, underparameterization and the Lucas critique,” *Journal of Monetary Economics*, 53(2), 249–264.
- FERRERO, G. (2007): “Monetary Policy, Learning and the Speed of Convergence,” *Journal of Economic Dynamics and Control*, 39(9), 3006–3041.
- GALI, J. (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction of the New Keynesian Framework*. Princeton University Press.
- GASPAR, V., F. SMETS, AND D. VESTIN (2006): “Adaptive Learning, Persistence and Optimal Monetary Policy,” *Journal of the European Economics Association*, 4, 376–385.
- (2010): “Inflation Expectations, Adaptive Learning and Optimal Monetary Policy,” in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and M. Woodford, vol. 3 of *Handbook of Monetary Economics*, chap. 19, pp. 1055–1095. Elsevier.
- GAUS, E. (2014): “Time-Varying Parameters and Endogenous Learning Algorithms,” unpublished, Ursinus College.
- GELAIN, P., K. J. LANSING, AND C. MENDICINO (2013): “House Prices, Credit Growth, and Excess Volatility: Implications for Monetary and Macroprudential Policy,” *International Journal of Central Banking*, 9(2), 219–276.
- GELAIN, P., K. J. LANSING, AND G. J. NATVIK (2015): “Explaining the boom-bust cycle in the U.S. housing market: a reverse-engineering approach,” Working Paper Series 2015-2, Federal Reserve Bank of San Francisco.
- GIANNONI, M. P., AND M. WOODFORD (2002): “Optimal Interest-Rate Rules: I. General Theory,” NBER Working Paper no. 9419.
- (2010): “Optimal Target Criteria for Stabilization Policy,” unpublished, Columbia University.
- GIBBS, C. (2015): “Forecast Combination, Non-linear Dynamics, and the Macroeconomy,” forthcoming, Economic Theory, UNSW.
- GIBBS, C., AND M. KULISH (2015): “Disinftions in a Model of Imperfectly Achored Expectations,” unpublished, UNSW.
- GRANDMONT, J.-M. (1998): “Expectations Formation and Stability of Large Socioeconomic Systems,” *Econometrica*, 66(4), 741–782.
- GÜRKAYNAK, R. S., B. SACK, AND E. SWANSON (2005): “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models,” *American Economic Review*, 95(1), 425–436.

- HONKAPOHJA, S., AND K. MITRA (2005): “Performance of Monetary Policy with Internal Central Bank Forecasting,” *Journal of Economic Dynamics and Control*, 29(4), 627–658.
- (2014): “Targeting nominal GDP or prices: Guidance and expectations dynamics,” Bank of Finland Research Discussion Papers.
- HOWITT, P. (1992): “Interest Rate Control and Nonconvergence to Rational Expectations,” *Journal of Political Economy*, 100(4), 776–800.
- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2013): “Is There a Trade-Off between Inflation and Output Stabilization?,” *American Economic Journal: Macroeconomics*, 5(2), 1–31.
- KITNEY, P. (2014): “Financial Factors and Monetary Policy: Determinacy and Learnability of Equilibrium,” unpublished manuscript.
- KOSTYSHYNA, O. (2012): “Application of an Adaptive Step-Size Algorithm in Models of Hyperinflation,” *Macroeconomic Dynamics*, 16, 355–375.
- KOZICKI, S., AND P. A. TINSLEY (2012): “Effective Use of Survey Information in Estimating the Evolution of Expected Inflation,” *Journal of Money, Credit and Banking*, 44(1), 145–169.
- 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.
- KRYVTSOV, O., M. SHUKAYEV, AND A. UEBERFELDT (2008): “Adopting Price-Level Targeting Under Imperfect Credibility: An Update,” Working Paper 2008-37, Bank of Canada.
- KUROZUMI, T., AND W. V. ZANDWEGHE (2014): “A Pitfall of Expectational Stability Analysis,” Federal Reserve Bank of Kansas RWP 14-07.
- KYDLAND, F. E., AND E. C. PRESCOTT (1982): “Time to Build and Aggregate Fluctuations,” *Econometrica*, 50(6), 1345–1370.
- LANSING, K. J. (2002): “Real-time estimation of trend output and the illusion of interest rate smoothing,” *Economic Review*, pp. 17–34.
- LAUBACH, T., AND J. C. WILLIAMS (2003): “Measuring the Natural Rate of Interest,” *The Review of Economics and Statistics*, 85(4), 1063–1070.
- LEEPER, E. (1991): “Equilibria Under ‘Active’ and ‘Passive’ Monetary and Fiscal Policies,” *Journal of Monetary Economics*, 27, 129–147.
- LEVIN, A., V. WIELAND, AND J. C. WILLIAMS (1999): “Robustness of Simple Monetary Policy Rules under Model Uncertainty,” in *Monetary Policy Rules*, ed. by J. Taylor. University of Chicago Press, Chicago.
- (2003): “Robustness of Forecast-Based Monetary Policy Rules under Model Uncertainty,” *American Economic Review*, 93, 622–645.

- LLOSA, L.-G., AND V. TUESTA (2008): “Determinacy and Learnability of Monetary Policy Rules in Small Open Economies,” *Journal of Money, Credit and Banking*, 40(5), 1033–1063.
- (2009): “Learning about monetary policy rules when the cost-channel matters,” *Journal of Economic Dynamics and Control*, 33(11), 1880–1896.
- LUBIK, T. A., AND C. MATTHES (2014): “Indeterminacy and Learning: An Analysis of Monetary Policy in the Great Inflation,” Federal Reserve Bank of Richmond WP 14-02.
- MANKIW, N. G., R. REIS, AND J. WOLFERS (2004): “Disagreement about Inflation Expectations,” in *NBER Macroeconomics Annual 2003, Volume 18*, NBER Chapters, pp. 209–270. National Bureau of Economic Research, Inc.
- MARCET, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93(5), 1476–1498.
- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models,” *Journal of Economic Theory*, (48), 337–368.
- MCCALLUM, B. T. (1981): “Price-level Determinacy With an Interest Rate Policy Rule and Rational Expectations,” *Journal of Monetary Economics*, 8, 319–329.
- (1983): “On Non-uniqueness in Rational Expectations Models: An Attempt at Perspective,” *Journal of Monetary Economics*, 11, 139–168.
- (1999): “Issues in the Design of Monetary Policy Rules,” in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford. North-Holland, Amsterdam.
- (2007): “E-stability vis-a-vis determinacy results for a broad class of linear rational expectations models,” *Journal of Economic Dynamics and Control*, 31(4), 1376–1391.
- MCGOUGH, B., G. D. RUDEBUSCH, AND J. C. WILLIAMS (2005): “Using a Long-term Interest Rate as the Monetary Policy Instrument,” *Journal of Monetary Economics*, (52), 855–879.
- MELE, A., K. MOLNAR, AND S. SANTORO (2012): “The Suboptimality of Commitment Equilibrium when Agents are Learning,” unpublished, Bank of Italy.
- MILANI, F. (2007): “Expectations, Learning and Macroeconomic Persistence,” *Journal of Monetary Economics*, 54, 2065–2082.
- (2008): “Learning, monetary policy rules, and macroeconomic stability,” *Journal of Economic Dynamics and Control*, 32(10), 3148–3165.
- (2011): “Expectation Shocks and Learning as Drivers of the Business Cycle,” *Economic Journal*, 121(552), 379–401.
- (2012): “The Modeling of Expectations in Empirical DSGE Models: a Survey,” Working Papers 121301, University of California-Irvine, Department of Economics.
- (2014): “Learning and time-varying macroeconomic volatility,” *Journal of Economic Dynamics and Control*, 47(C), 94–114.

- MOLNAR, K., AND S. SANTORO (2013): “Optimal Monetary Policy When Agents Are Learning,” *European Economic Review*, 66, 39–62.
- ORMENO, A., AND K. MOLNAR (2013): “Using Survey Data of Inflation Expectations in the Estimation of Learning and Rational Expectations Models,” manuscript, Norwegian School of Economics.
- ORPHANIDES, A. (2001): “Monetary Policy Rules Based on Real-Time Data,” *American Economic Review*, 91(4), 964–985.
- ORPHANIDES, A., AND J. C. WILLIAMS (2002): “Robust Monetary Policy Rules with Unknown Natural Rates,” *Brookings Papers on Economic Activity*, 33(2), 63–146.
- (2005a): “The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations,” *Journal of Economic Dynamics and Control*, 29(11), 1927–1950.
- (2005b): “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in *The Inflation Targeting Debate*, ed. by B. S. Bernanke, and M. Woodford. University of Chicago Press.
- (2005c): “Inflation scares and forecast-based monetary policy,” *Review of Economic Dynamics*, 8(2), 498–527.
- (2007): “Robust monetary policy with imperfect knowledge,” *Journal of Monetary Economics*, 54(5), 1406–1435.
- (2008): “Learning, expectations formation, and the pitfalls of optimal control monetary policy,” *Journal of Monetary Economics*, 55(Supplemen), S80–S96.
- (2012): “Monetary Policy Mistakes and the Evolution of Inflation Expectations,” in *The Great Inflation: The Rebirth of Modern Central Banking*, NBER Chapters, pp. 255–288. National Bureau of Economic Research, Inc.
- PINTUS, P. A., AND J. SUDA (2014): “Learning Financial Shocks and the Great Recession,” Working papers, HAL.
- 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.
- (2008): “Adaptive Learning and the Use of Forecasts in Monetary Policy,” *Journal of Economic Dynamics and Control*, 32(4), 2661–3681.
- PRIMICERI, G. E. (2005): “Why Inflation Rose and Fell: Policymakers’ Beliefs and US Postwar Stabilization Policy,” *The Quarterly Journal of Economics*, 121, 867–901.
- REIFSCHNEIDER, D., AND J. C. WILLIAMS (2000): “Three lessons for monetary policy in a low-inflation era,” *Journal of Money, Credit and Banking*, 32(4), 936–978.
- ROBERTS, J. M. (1995): “New Keynesian Economics and the Phillips Curve,” *Journal of Money, Credit and Banking*, 27, 975–984.

- ROTEMBERG, J. J., AND M. WOODFORD (1999): “Interest-Rate Rules and an Estimated Sticky-Price Model,” in *Monetary Policy Rules*, ed. by J. Taylor, pp. 57–126. University of Chicago Press, Chicago.
- SARGENT, T., N. WILLIAMS, AND T. ZHA (2006): “Shocks and Government Beliefs: The Rise and Fall of American Inflation,” *American Economic Review*, 96(4), 1193–1224.
- 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.
- SINHA, A. (2014): “Learning and the Yield Curve,” manuscript, Santa Clara University.
- (2015): “Government debt, learning and the term structure,” *Journal of Economic Dynamics and Control*, 53(C), 268–289.
- SLOBODYAN, S., AND R. WOUTERS (2012a): “Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models,” *American Economic Journal: Macroeconomics*, 4(2), 65–101.
- (2012b): “Learning in an estimated medium-scale DSGE model,” *Journal of Economic Dynamics and Control*, 36(1), 26–46.
- SMETS, F., AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: a Bayesian DSGE Approach,” *American Economic Review*, 97(3), 586–606.
- STOCK, J. H., AND M. W. WATSON (1989): “New Indexes of Coincident and Leading Economic Indicators,” in *NBER Macroeconomics Annual 1989, Volume 4*, NBER Chapters, pp. 351–409. National Bureau of Economic Research, Inc.
- (2007): “Why Has U.S. Inflation Become Harder to Forecast?,” *Journal of Money, Credit and Banking*, 39(s1), 3–33.
- SVENSSON, L. E. (2003): “What Is Wrong with Taylor Rules? Using Judgement in Monetary Policy through Targeting Rules,” *Journal of Economic Literature*, 41, 426–477.
- TAYLOR, J. (1993): “Discretion Versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.
- TETLOW, R. J., AND P. VON ZUR MUEHLEN (2001): “Simplicity versus optimality: The choice of monetary policy rules when agents must learn,” *Journal of Economic Dynamics and Control*, 25(1-2), 245–279.
- WILLIAMS, J. C. (2010): “Monetary policy in a low inflation economy with learning,” *Economic Review*, pp. 1–12.
- WILLIAMS, N. (2014): “Escape Dynamics in Learning Models,” unpublished, University of Wisconsin - Madison.

- WINKLER, F. (2015): “The Role of Learning for Asset Prices, Business Cycles and monetary Policy,” manuscript, London School of Economics.
- WOODFORD, M. (1998): “Public Debt and the Price Level,” unpublished manuscript, Columbia University.
- (2001): “Fiscal Requirements of Price Stability,” *Journal of Money, Credit and Banking*, 33, 669–728.
- (2003a): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- (2003b): “Optimal Interest-Rate Smoothing,” *Review of Economic Studies*, 70, 861–886.
- (2005): “Comments on: “Using a Long-term Interest Rate as the Monetary Policy Instrument”,” *Journal of Monetary Economics*, (52), 881–887.
- (2010): “Robustly Optimal Monetary Policy with Near-Rational Expectations,” *American Economic Review*, 100(1), 274–303.
- (2013): “Macroeconomic Analysis Without the Rational Expectations Hypothesis,” *Annual Review of Economics*, 5(1), 303–346.

LIST OF FIGURES

Figure 1: The figure shows the evolution of survey forecast for the GDP deflator. In particular, the light grey dotted (dark grey dashed) line shows 1-quarter (4-quarters) ahead consensus forecast from the Survey of Professional Forecasters. The black line with circles shows the 5-10 years ahead consensus forecast from Blue Chip Financial Forecasts survey.

Figure 2: The figure shows the evolution of survey forecast for the 3-months T-bill interest rate. In particular, the light grey dotted (dark grey dashed) line shows 1-quarter (4-quarters) ahead consensus forecast from the Survey of Professional Forecasters. The black line with circles shows the 5-10 years ahead consensus forecast from Blue Chip Financial Forecasts survey.

Figure 3: The figure shows the evolution of survey forecast for the real GDP growth. In particular, the light grey dotted (dark grey dashed) line shows 1-quarter (4-quarters) ahead consensus forecast from the Survey of Professional Forecasters. The black line with circles shows the 5-10 years ahead consensus forecast from Blue Chip Financial Forecasts survey.

Figure 4: The figure shows the (weighted) volatilities out inflation and output gap associated to alternative targeting rules. The black dashed line corresponds to targeting the model with learning, while the grey solid line show pertains the model under rational expectations.

Figure 5: The figure shows the response to a 1-standard deviation positive cost push shock. In all panels, the dashed black line corresponds to optimal commitment, the solid black line shows optimal policy under learning and the solid grey line depict the response of the economy under the targeting rule discussed in section 6.1.

Figure 6: The figure shows the volatility of output and interest rate as a function of the constant gain $\bar{g}$. The black lines display the volatility output gap (triangles) and interest rate (circles) in the case where the central bank associates no cost to interest rate volatility. The grey dashed (solid) line show the volatility of output (interest rate) when the policymaker cares about interest rate volatility.

Figure 7: The figure shows the (weighted) volatilities of inflation and output gap associated to alternative policies as a function of λ_x . In particular, the black solid line shows the volatility under the optimal policy under learning. The dashed grey line shows the volatility for optimal policy under learning when the central bank is in full control of aggregate demand. Finally the solid grey line shows optimal discretion.

Figure 8: The figure shows stability frontiers corresponding to alternative Taylor rules. In particular $(\bar{g}, \phi_\pi)$ pairs above the frontier correspond to locally unstable equilibria under constant gain learning. The black solid line corresponds to the standard Taylor specification with $\phi_x^* = 0.5/4$. The solid (dashed) grey line corresponds to $\phi_x = \phi_x^*/2$ ($\phi_x = \phi_x^*/3$).

Figure 9: The figure shows the phase diagram for inflation and the output gap in the model where the monetary policy rule is not communicated to market participants.

Figure 10: The figure shows the phase diagram for inflation and the output gap in the model where agents know the monetary policy rule.