

Monetary Policy Design under Imperfect Knowledge: An Open Economy Analysis

Yu-chin Chen ^a and Pisut Kulthanavit ^{b, *}

First draft: May 2008

This draft: May 2009

Abstract. Incorporating adaptive learning into an open-economy DSGE model, we examine how monetary policy rules should adjust when agents' information set deviates from that assumed under the rational expectations paradigm. We find that when agents observe current shocks but don't know the parameters governing key macroeconomic dynamics, the resulting distortion is small and the preferred policy under rational expectations works well. However, the welfare cost of imperfect knowledge becomes quite severe when agents also have to learn about the structural shocks to the economy. Monetary policy can play a significant role in mitigating distortions associated with this form of imperfect knowledge.

JEL classification: D84; E52; F41

Keywords: Adaptive learning; Monetary policy rules; Open economy

We thank, without implicating, Richard Startz and Stephen Turnovsky for useful comments and suggestions. We also benefited from seminar participants at the University of Washington, Hong Kong University, 2009 WEAI Pacific Rim conference, 2009 LACEA conference, and the 2007 Learning Week conference at the St. Louis Federal Reserve Bank. Any remaining errors are our own.

^a Department of Economics, 206 Condon Hall, Box 353330, University of Washington, Seattle, WA, 98195, U.S.A.; Tel: 1-206-543-6197; Fax: 1-206- 685-7477; *E-mail:* yuchin@u.washington.edu

^b Department of Economics, 2 Phrachan Road, Phranakorn, Thammasat University, Bangkok, 10200, Thailand.; Tel: 66-2613-2471; Fax: 66-2224-9428; *E-mail:* pisut@econ.tu.ac.th

* *Corresponding Author*

1. Introduction

How should policymakers in an open economy respond to inflationary pressure originating at home versus that originating abroad, while maintaining domestic output targets? Recent studies including Clarida, Gali and Gertler (2001, 2002) and Gali and Monacelli (2005) provide comprehensive answers to this important question using the New-Keynesian dynamic stochastic general equilibrium (DSGE) framework. However, since these studies generally impose the rational expectations (RE) assumption, they cannot assess the effects of information imperfection on the outcomes of alternative policies. As structural shifts are noticeably common in many economies, it is important to know whether a chosen policy under RE would still work well when agents do not know the structural dynamics or the sources of shocks in the economy.

In this paper, we specifically tackle the issue of how monetary policy rules should adjust when agents' information sets deviate from those assumed under the RE paradigm. By incorporating adaptive learning into a standard open economy DSGE model, we analyze the conditions under which policymakers should target *domestic producer price* inflation (DI) versus *consumer price* inflation (CI).¹ In particular, we consider situations in which the dominant source of economic shocks - domestic versus foreign - dictates certain policy prescription under RE; we then examine how the degree of knowledge imperfection and the associated learning behavior affect welfare and the preferred policy choice. We find that not all forms of ignorance are bliss; instead, the welfare cost can be quite severe when agents do not know the contemporaneous structural disturbances affecting the economy. In such instances, monetary policy should deviate from the preferred choice under RE in order to help lessen the cost of learning.

The RE framework has dominated macroeconomic modeling and policy analyses since the 1970s, but the assumption that people's expectations coincide with the forecasts generated by the models that describe their behavior is not always appropriate. For example, Sargent (1993) states, "rational expectations models impute much more knowledge to the agents within the model than is possessed by an econometrician, who faces estimation and inference problems." In addition, numerous recent studies have emphasized the prevalence of structural breaks in key

¹ In our main analyses, we assume the central bank to follow a standard forward-looking Taylor rule with fixed weights on the output gap and a choice of either CI or DI.

macroeconomic dynamics, including the size and the volatility of output shocks, the degree of exchange rate pass-through, and the persistence of inflation dynamics.² These underlying instabilities, which are not easy to detect and predict even for econometricians, call into question the plausibility that economic agents always know or can forecast correctly the RE equilibrium. To explore how deviations from RE and the associated forecast errors could affect welfare and policy choices, we model economic agents, including policymakers, as bounded-rational, and they rely on an adaptive learning mechanism, à la Evans and Honkapojah (2001), to form expectations.³ The adaptive learning approach has become increasingly popular as a tractable modeling alternative to rational expectations, and recently, a growing number of central bank research has applied it to monetary policy analyses (see, for example, Orphanides and Williams 2005, 2007a,b).⁴ Previous research focuses mostly on closed-economy policy analyses, and does not specifically address the interplay between the extent of knowledge imperfection and the sources of dominant economic disturbances, which we believe to be an issue of practical importance for the conduct of monetary policy in a small open economy.

We set up an open economy version of the standard DSGE model with nominal price rigidity, similar to that in Gali and Monacelli (hereafter, GM 2005). The economy is subject to three types of structural shocks: domestic productivity shock, foreign inflation shock, and a domestic cost-push shock that helps capture the policy tradeoff in stabilizing inflation versus output. The policymaker is assumed to follow an instrument rule and chooses between two variants of the forward-looking Taylor rule, where the interest rate adjusts in response to forecasts of the output gap and the deviations of inflation from its target. In one case, the policymaker sets a DI target, and we call this rule “F-DITR” for forward-looking DI Taylor rule. In the other, a CI target is chosen, so we call it a forward-looking CI Taylor rule (“F-CITR”). The main difference between the two rules is that F-CITR indirectly responds to foreign shocks, which influence the economy

² For example, Kim and Nelson (1999), Stock and Watson (2003) amongst others, document the “great moderation” of output growth and inflation volatilities since the mid-1980s, and Cogley, Primiceri, and Sargent (2008) shows evidence for a time-varying inflation target for the U.S.

³ This paper focuses on the welfare evaluation of monetary policy instead of expectational stability (E-Stability) of an RE equilibrium. Bullard and Schaling (2006) and Llosa and Tuesta (2008), for example, study E-Stability of various policy rules in an open economy. Evans and Honkapojah (2001) discusses E-Stability in more details.

⁴ One could consider the learning behavior as transitional towards the new RE steady state after a structural break has occurred, or as a more realistic description of expectation formation in general. To better capture real time learning, we assume private agents engage in constant-gain or “perpetual learning,” where they put more weight on newly available data in the learning process to account for the possibility of structural changes.

through the terms of trade. While these simple rules may not be the optimal welfare-maximizing choice, we believe they are more realistic and practically relevant due to their implementability.⁵ In addition, these simple rules also provide an easy platform for us to focus more effectively on the impact of imperfect knowledge. Assuming away RE has a number of consequences; the chief one of which is that apart from its stabilizing role, monetary policy may further facilitate the agents' learning process and mitigate the adverse effects of imperfect knowledge. Our DSGE setup allows explicit welfare calculations based on the simulated dynamics of output and inflation as they respond to alternative policy rules.

We consider two different degrees of knowledge imperfection in attempt to capture likely real-world scenarios in the expectation formation processes. Relaxing the RE assumption that agents have full knowledge of the underlying economic dynamics, we first assume that they know the functional form but not the parameter values associated with the rational expectations equilibrium. This scenario, in our view, mimics the information set of agents in an economy that has either just experienced or is subject to possible regime shifts. It is also the assumption in recent literature such as in Bullard and Duffy (2004), Chakraborty and Evans (2008) and Mark (2009). Next, we consider a more severe form of knowledge imperfection under which the agents also lack information about the current structural shocks to the economy. (This form of knowledge imperfection is assumed in Orphanides and Williams (2005, 2007a).) Juxtaposed to the first scenario, this setup allows us to assess the potential benefit of releasing and announcing available new data in a speedy manner. In both scenarios, agents rely on observable historical data and an adaptive learning mechanism to make estimates and form expectations. The associated forecast errors from this learning process can then affect policy outcomes and welfare evaluations.⁶

Our simulation results under RE first confirm the conventional wisdom that the relative volatility of domestic versus foreign shocks can alter the preferred policy choice. When foreign

⁵ Our robustness checks show that the qualitative results hold across a reasonable set of alternative fixed weights for the interest rate response to output gap and inflation deviations, and that the additional welfare gain from fine-tuning the weights are small in magnitude compare to the results we focus on.

⁶ Using a closed-economy learning model, Orphanides and Williams (2005, 2007a) point out that under imperfect knowledge, a policy rule may need to be more aggressive on inflation control, in order to anchor inflation expectations and facilitate learning. Preston (2005, 2008) studies monetary policy designs with an alternative approach of adaptive learning in which long-horizon expectations of private agents influence their consumption and production decisions rather than one-period-ahead expectations.

shocks are relatively quiet, targeting CI is preferable. Reaction to foreign shocks in F-CITR brings them into the domestic macro dynamics. When foreign shocks are more stable, so is the terms of trade, and the induced expenditure-switching impact on domestic output is less pronounced. Despite the increase in output volatility induced from this channel, anchoring on the stable foreign nominal shocks helps reduce domestic inflation volatility considerably. Since the welfare loss function penalizes inflation variation more heavily relative to output volatility, F-CITR policy rule delivers higher welfare through the nominal anchor effect (or more precisely, lower welfare loss from the first best world). On the other hand, when foreign inflation shocks are very volatile relative to the domestic shocks, the policy rule that targets DI significantly dominates F-CITR in welfare evaluations. In sum, under RE, policymakers should put a larger weight or response elasticity on the more stable target. This conclusion confirms results based on classical theories as well as other recent studies.⁷

Turning to adaptive learning, our simulations under the two cases of imperfect knowledge result in the following findings. First, relaxing the assumption that economic agents have full information concerning the exact dynamic equations describing the RE equilibrium *alone* does not alter the preferred policy outcome significantly. That is, when agents can observe current shocks but have to learn the parameter values governing the relevant economic dynamics, the excess volatility induced by the learning process and the associated welfare losses are mostly negligible. Policymakers can relatively safely follow the policy prescription derived under the RE analyses, and target away from the dominant, more volatile shocks. However, if agents do not have information about the contemporaneous structural shocks hitting the economy and have to learn their dynamics in addition, the induced welfare consequences become much more severe, and monetary policy needs to play a more active role in helping the agents learn. In particular, regardless of the source of the dominant shocks, we find a DI targeting rule to be preferable. Under limited information and more uncertainty, targeting fewer variables helps provide a sharper anchor to facilitate learning. The domestic inflation targeting rule does not respond to the unobserved foreign shocks, and thus dampens forecast errors and improves welfare. This finding suggests that public dissemination of up-to-date economic information can significantly

⁷ This finding can also be justified in a ‘signal extraction’ framework. In a multi-sector model, Mankiw and Reis (2003) show that policymakers should put more weight on stabilizing the prices of sectors that are subject to less idiosyncratic shocks.

improve welfare, and in the case where it is not possible, monetary policy rule should adjust from the RE strategy in order to lessen the welfare loss induced by the learning process.

The rest of the paper is organized as follows. Section 2 outlines the open economy general equilibrium model and discusses the monetary policy rules under examination. Section 3 presents the equilibrium concepts and solution methodology for rational expectations and adaptive learning. Section 4 discusses the calibration and simulation procedures and presents our findings. Finally, we conclude in Section 5.

2. The Small Open Economy Model and Monetary Policy Rules

One salient policy debate on monetary policy design in an open economy is whether, and under what conditions, policymakers should target *domestic producer price* inflation (DI) versus *consumer price* inflation (CI). This question is related to the broader literature on how to set inflation targets in a multi-goods environment (recent examples include Aoki (2001) and Mankiw and Reis (2003)), yet the open economy aspect has its obvious practical policy relevance, and thus draws a significant amount of research interest on its own. Results based on recent open-economy DSGE models such as Clarida et al (2001, 2002) and Gali and Monacelli (2005) point to domestic price stabilization. Their benchmark framework, with producer currency price rigidity, complete asset markets, and frictionless trade and labor market, assumes commitment mechanism and the availability of fiscal policy to counteract other forms of distortions in the economy. Under such a setting, the role of monetary policy is to undo domestic price rigidity, which is the only distortion left to correct in the economy.⁸ Subsequent studies examine alternative setups such as local currency pricing, the presence of labor market frictions and a non-tradable sector, and as a result, CPI-targeting or even wage targeting may come to dominate DI stabilization (see, for example, Corsetti and Pesenti (2005), Campolmi

⁸ It is common in literature to assume an employment subsidy that offsets the distortion from monopolistic competition. In a closed economy setting, the only distortion remaining is from nominal price rigidity. Thus the goal of monetary policy is to correct this distortion by stabilizing markups of all domestic firms at their flexible price level (see Gali (2003)). In an open economy, Corsetti and Pesenti (2001) and Benigno and Benigno (2003) note that there can be an additional terms of trade distortion. Under a Nash equilibrium, each monetary authority would have an incentive to manipulate the terms of trade to improve its own country's welfare. Under this scenario, the employment subsidy that offsets the monopolistic power distortion is not sufficient to render the flexible price equilibrium allocation optimal. GM (2005) shows under certain parameter assumption, one can ignore the terms of trade consideration and restore the goal of monetary policy as to address the price rigidity distortion only.

(2008), Doyle, Erceg, and Levin (2006) and others). While the exact transmission mechanism may be model-specific, one general lesson from this literature is that monetary policy may have to deviate from the benchmark result of replicating the flexible price equilibrium when the economy faces additional frictions, such as labor market rigidity or the knowledge imperfection we consider. Below we adopt a setup similar to GM (2005) and abstract away from other forms of frictions discussed in the broader literature, in order to focus on the effect of deviations from RE in the form of adaptive learning.

2.1 The Model

Our baseline general equilibrium model consists of a continuum of identical small open economies uniformly distributed on the unit interval, as in GM (2005). Monopolistically competing producers use production technology linear in labor input and face Calvo (1983) staggered price-setting. Securities markets are assumed to be complete, and *purchasing power parity* (PPP) holds. We close the dynamic system with alternative monetary rules, expressed as the forward-looking specification of Taylor-typed rules. As preferences, production technology, and market structures are assumed to be symmetric across countries, below we present the optimization problems facing the representative household and firm from the perspective of one of these economies, indexed by H (*Home*). We treat the rest of the *world* as a foreign block, with corresponding variables denoted by a superscript "*" ⁹.

Households

Each economy is populated by a representative consumer who maximizes expected discounted utility from consumption and labor-leisure choice. The representative household maximizes the following utility function:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} - \frac{N_t^{1+\varphi}}{1+\varphi} \right]$$

where β is the household discount factor, σ is the inverse of the intertemporal elasticity of consumption, and φ is the inverse of labor supply elasticity. N_t denotes the labor supply.

Consumption index C_t is a CES composite of home and foreign good consumption, defined by:

⁹ See GM (2005) for a more detailed discussion of this benchmark setup.

$$C_t \equiv \left[(1-\alpha)^{1/\eta} (C_{H,t})^{(\eta-1)/\eta} + \alpha^{1/\eta} (C_{F,t})^{(\eta-1)/\eta} \right]^{\eta/(\eta-1)}$$

where $\eta > 0$ measures the elasticity of substitution between domestic and foreign goods, and $C_{H,t}$ and $C_{F,t}$ each are CES aggregated consumption indices of home and imported goods, with the elasticity of substitution among goods within each category given by ε and γ . $\alpha \in [0,1]$ represents the share of domestic consumption allocated to imported goods and can be interpreted as a degree of trade openness.

The sequence of household's budget constraints takes the form:

$$\int_0^1 P_{H,t}(j)C_{H,t}(j)dj + \int_0^1 \int_0^1 P_{i,t}(j)C_{i,t}(j)djdi + E_t [Q_{t,t+1}D_{t+1}] \leq D_t + W_t N_t + T_t$$

for $t = 0, 1, 2, \dots$, where $P_{H,t}(j)$ and $C_{H,t}(j)$ are the price and consumption of the j th domestic goods respectively, and $P_{i,t}(j)$ and $C_{i,t}(j)$ are the price and consumption of the j th variety imported from country i in terms of domestic currency. $Q_{t,t+1}$ denotes the stochastic discount factor for one-period ahead nominal pay-offs and D_{t+1} is the nominal pay-off in period $t + 1$ of the household's portfolio at the end of period t . W_t is the nominal wage and T_t is lump-sum transfers/taxes.

The consumer price index (CPI) takes the form:

$$P_t \equiv \left[(1-\alpha)(P_{H,t})^{1-\eta} + \alpha(P_{F,t})^{1-\eta} \right]^{1/(1-\eta)}$$

where $P_{H,t}$ and $P_{F,t}$ each are CES aggregated price indices of domestic and imported goods.

Domestic Producers

On the production side, we assume that monopolistically-competing firms produce with technology linear in labor and set prices in a staggered fashion à la Calvo (1983). We let parameter θ denote the fraction of firms per period that cannot adjust their prices, so at each time t , a fraction $1 - \theta$ of firms set new prices optimally. The firm j sets the new price $\bar{P}_{H,t}$ in period t to maximize the expected discounted value of profits:

$$E_t \sum_{k=0}^{\infty} \theta^k \left\{ Q_{t,t+k} \left[Y_{t+k}(j) (\bar{P}_{H,t} - MC_{t+k}^n) \right] \right\}$$

subject to the sequence of the demand curves:

$$Y_{t+k}(j) \leq \left(\frac{\bar{P}_{H,t}}{P_{H,t+k}} \right)^{-\varepsilon} \left(C_{H,t+k} + \int_0^1 C_{H,t+k}^i di \right)$$

where $Y_t(j)$ is the output of firm j and MC_t^n is the nominal marginal cost.

Equilibrium

Solving for the market clearing conditions, our small open economy is described by the following log-linearized equilibrium dynamics, as in GM (2005). The first equation is a forward-looking IS equation from the clearing of the goods market:

$$x_t = E_t x_{t+1} - \frac{1}{\sigma_\alpha} \left(r_t - E_t \pi_{H,t+1} - \bar{r}r_t \right) \quad (1).$$

Here x_t and r_t denote the output gap and the domestic interest rate, respectively, and $\bar{r}r_t$ is the domestic natural rate of interest.¹⁰ r_t also represents the policy instrument which is endogenously set by the policymakers in the model. $\pi_{H,t} = p_{H,t} - p_{H,t-1}$ is *domestic producer price* inflation (DI), where $p_{H,t}$ is the (log) *domestic producer price* index. The home natural rate of interest, $\bar{r}r_t$, depends on the expected growth rate of world output, and (log) labor productivity v_t .¹¹ We assume that v_t follows an AR(1) process $v_t = \rho_v v_{t-1} + \varepsilon_{v,t}$.

The next equilibrium condition is the New-Keynesian Phillips curve (NKPC):

$$\pi_{H,t} = \beta E_t \pi_{H,t+1} + \kappa_\alpha x_t + u_t \quad (2)$$

where κ_α , the slope coefficient, depends on the degree of openness and other parameters and is defined in footnote 10. A cost-push shock, u_t , is introduced to capture a non-trivial tradeoff in stabilizing inflation and output that the policymakers face when they design a policy rule. We note that when this small open economy is in perfect autarky ($\alpha = 0$), the dynamic equations (1) and (2) are identical to the dynamic IS and NKPC equations, respectively, in a standard closed economy setup.¹²

¹⁰ For ease of presentation, we define the following new parameters in terms of the structural ones defined earlier: $\lambda \equiv [(1 - \beta\theta)(1 - \theta) / \theta]$, $\omega \equiv \sigma\gamma + (1 - \alpha)(\sigma\eta - 1)$, $\sigma_\alpha \equiv \sigma / [1 - \alpha + \alpha\omega]$, and $\kappa_\alpha \equiv \lambda(\sigma_\alpha + \varphi)$.

¹¹ $\bar{r}r_t \equiv \rho - \sigma_\alpha \Gamma(1 - \rho_v)v_t + \alpha\sigma_\alpha(\Theta + \Psi)E_t[\Delta y_{t+1}^*]$, $\Gamma \equiv (1 + \varphi) / (\sigma_\alpha + \varphi)$, $\Theta \equiv \omega - 1$, and $\Psi \equiv -(\Theta\sigma_\alpha) / (\sigma_\alpha + \varphi)$. Δy_{t+1}^* is the growth rate of world output.

¹² See, for example, Clarida et al. (1999) and Woodford (2003).

We further assume that *purchasing power parity* holds, so the relationship between CI, $\pi_t = p_t - p_{t-1}$, and DI, $\pi_{H,t}$, is given by:

$$\pi_t = \pi_{H,t} + \alpha \Delta s_t \quad (3),$$

where $s_t \equiv p_{F,t} - p_{H,t}$ is the (log) effective terms of trade, $p_{F,t}$ is the (log) price index for imported goods (expressed in domestic currency), and p_t is the (log) consumer price index.¹³

Using this relationship, we can express the previous two equilibrium conditions (1) and (2) in terms of CI:

$$x_t = E_t x_{t+1} - \frac{1}{\sigma_\alpha} (r_t - E_t \pi_{t+1} - \bar{r}) - \frac{1}{\sigma_\alpha} \alpha E_t s_{t+1} + \frac{1}{\sigma_\alpha} \alpha s_t \quad (4),$$

$$\pi_t = \beta E_t \pi_{t+1} + \kappa_\alpha x_t - \alpha \beta E_t s_{t+1} + \alpha (1 + \beta) s_t - \alpha s_{t-1} + u_t \quad (5).$$

To describe the dynamics of the terms of trade, s_t , we note that under the assumption of complete international asset markets, *uncovered interest parity* (UIP) condition is expressed as the following:

$$r_t - r_t^* = E_t [\Delta e_{t+1}] \quad (6)$$

where e_t is the (log) nominal effective exchange rate and r_t^* is the *world* interest rate. Assuming that the *law of one price* holds for each individual good, we have $p_{F,t} = e_t + p_t^*$ where p_t^* is the (log) *world* price index. The (log) effective terms of trade would then be $s_t = e_t + p_t^* - p_{H,t}$. This expression also implies that:

$$\Delta s_t = \Delta e_t + \pi_t^* - \pi_{H,t} \quad (7)$$

where $\pi_t^* = p_t^* - p_{t-1}^*$ is world inflation. Combining (7) with the UIP condition (6), we obtain:

$$s_t = E_t s_{t+1} - (r_t - E_t \pi_{H,t+1}) + (r_t^* - E_t \pi_{t+1}^*) \quad (8).$$

Plugging (3) into (8), we obtain the following stochastic difference equation:

$$(1 - \alpha) s_t = (1 - \alpha) E_t s_{t+1} - (r_t - E_t \pi_{t+1}) + (r_t^* - E_t \pi_{t+1}^*) \quad (9).$$

where π_t^* is assumed to follow an AR(1) process $\pi_t^* = \rho_{\pi^*} \pi_{t-1}^* + \varepsilon_{\pi^*t}$.

¹³ Note that when the economy is completely autarkic, CI collapses to DI and the open economy model becomes identical to the closed economy counterpart.

Following GM (2005), we approximate the representative household's expected utility locally around the flexible price steady state using a second-order Taylor expansion, which gives us a measure of the utility loss relative to the first-best optimal policy.¹⁴ This loss is measured as the fraction of the steady-state level of consumption:

$$W = -\frac{(1-\alpha)}{2} \sum_{t=0}^{\infty} \beta^t \left[\frac{\varepsilon}{\lambda} \pi_{H,t}^2 + (1+\varphi)x_t^2 \right] \quad (10).$$

where λ is defined as $(1-\beta\theta)(1-\theta)/\theta$ as in footnote 10. Taking unconditional expectations and letting the discount factor approach unity, the expected welfare losses of any policy rule that deviates from the optimal one can be expressed in terms of the variance of DI and the variance of the output gap, as follows:

$$EW = -\frac{(1-\alpha)}{2} \left[\frac{\varepsilon}{\lambda} \text{var}(\pi_{H,t}) + (1+\varphi) \text{var}(x_t) \right] \quad (11).$$

We use equation (11) to evaluate the performances of alternative monetary policy rules which we discuss next.

2.2 Monetary Policy Rules

We assume that the policymakers implement a simple instrument rule. We specify monetary policy as a choice of two variants of the Taylor-typed rule in which the monetary authority sets the domestic interest rate (r_t) in response to one-period-ahead forecasts of the output gap and the deviation of inflation from its target. In our framework, the central bank does not have informational advantage; the policymaker and the private agents have common one-period-ahead expectations about the macroeconomic conditions.¹⁵

The first policy rule has the policymaker stabilizing DI and the output gap, which we call a forward-looking DI Taylor-typed rule (F-DITR). This policy rule is expressed as:

$$r_t = \rho + \pi_H^T + \varphi_{\pi_H} (E_t \pi_{H,t+1} - \pi_H^T) + \varphi_x E_t x_{t+1} \quad (12)$$

where π_H^T is the target DI level. $\varphi_{\pi_H} > 0$ and $\varphi_x > 0$ measure how aggressive the policymaker responds to any deviation of DI and the output gap from their target values, e.g., π_H^T and zero,

¹⁴ See GM (2005) for more detailed derivations.

¹⁵ Another way to interpret this common expectations assumption is that the policymakers use private expectations in setting the policy interest rate.

respectively. Parameter $\rho = \beta^{-1} - 1$ is the time discount rate and can be interpreted as a quarterly risk-less return in the steady state.

The second policy is a forward-looking CI Taylor-typed rule (F-CITR). Under this policy rule, the policymaker adjusts the interest rate r_t in response to one-period-ahead expectations of the output gap and CI, rather than DI:

$$r_t = \rho + \pi^T + \varphi_\pi (E_t \pi_{t+1} - \pi^T) + \varphi_x E_t x_{t+1} \quad (13)$$

where π^T is a target of CI. Parameter $\varphi_\pi > 0$ measures the aggressiveness of policymakers to any deviation of CI from its target value (π^T).

The key difference between the above two policy rules is that F-CITR indirectly responds to foreign shocks which influence the domestic economy via the terms of trade, but F-DITR does not. For simplicity, we assume the world policymaker implements the a simple forward-looking Taylor rule as follows: $r_t^* = \rho + \pi^{T*} + \varphi_{\pi^*} (E_t \pi_{t+1}^* - \pi^{T*})$, where π^{T*} is a target of world inflation, which we set to zero.

3. Models of Expectations Formation

Private sector expectations play a crucial role in monetary policy implementation. The traditional RE paradigm presupposes that private agents possess full information, and as such their expectations coincide with what the model produces on average. This assumption may be too strong. As argued in Sargent (1993), private agents may be more appropriately modeled as bounded rational and possessing only partial knowledge about the economy when they form expectations. Below we present two forms of such knowledge imperfection and the adaptive learning framework that we use to model their expectation formation process. We start with a brief discussion of the benchmark RE setup to help illustrate the differences.

3.1 Rational Expectations (RE)

Under RE, private agents have perfect knowledge about the structure of the economy and efficiently use the information to form expectations. In other words, private agents know the rational expectations equilibrium (REE).

To solve the REE, we combine the dynamical system given by (1), (2) and (8) with the monetary policy rule (12), when the policymakers adopt F-DITR. For F-CITR, we complement the system given by (4), (5) and (9) with the policy rule (13). The reduced form system can be written as:

$$y_t = A + BE_t y_{t+1} + Cy_{t-1} + Dw_t \quad (14)$$

$$w_t = \rho_w w_{t-1} + \varepsilon_t \quad (15)$$

where $y_t = [x_t, \pi_{H,t}, s_t]'$, $w_t = [v_t, u_t, \pi_t^*]'$, and $\varepsilon_t = [\varepsilon_{v,t}, \varepsilon_{u,t}, \varepsilon_{\pi^*,t}]'$ with appropriate matrices A, B, C, and D.¹⁶ The *minimum state variable* (MSV) solution to the system given by equations (14) and (15) takes the form:

$$y_t = \bar{a} + \bar{b}y_{t-1} + \bar{c}w_t \quad (16)$$

where \bar{a} , \bar{b} and \bar{c} are conformable matrices.¹⁷ In sum, under RE, agents know the correct functional form (16) and its relevant parameter values in matrices \bar{a} , \bar{b} and \bar{c} . Agents then make uses of this knowledge to form their expectations.

3.2 Perpetual Learning

We model the adaptive learning process in this paper following the framework proposed by Evans and Honkapojah (2001) and Orphanides and Williams (2005, 2007a, b). In contrast to RE, private agents are bounded rational in that they only know the functional form of REE, and rely on an adaptive learning process to obtain relevant parameter estimates in order to form expectations. Under certain conditions where there is expectational stability (or stability under learning), forecast errors are corrected gradually over time and the economy will converge to the desired REE.

We model private agents as engaging in ‘perpetual learning’ or ‘constant gain’ learning. With this form of learning, we implicitly assume that agents are constantly attentive to the possibility of structural changes in the economy, so in estimating the REE, they pay more

¹⁶ For F-CITR, $y_t = [x_t, \pi_t, s_t]'$.

¹⁷ The MSV solution is generally considered a unique solution that is free of bubble and sunspot components. See McCallum (1983, 1998).

attention to recent data than historical information. As discussed in Orphanides and Williams (2007a, b), this form of adaptive learning is especially appropriate in economies undergoing structural or policy shifts, where agents are uncertain about the stability of the structural forces driving the economic dynamics. Alternatively, we can also interpret the fixed gain constant as a reflection of the degree of rationality. Since least squares learning, where the gain is declining inversely with sample size, eventually “converges” to REE as time horizon increases. We can thus interpret a smaller fixed gain as indicative of a higher degree of rationality. In addition, as discussed in Waters (2009), during periods of credible policy and economic stability, the public can rely on a longer range of historical data to learn about the structural parameters, rather than focusing more on the most recent news. This is also captured by a smaller gain constant.

The fundamental element of adaptive learning is that at each time t , private agents have a *Perceived Law of Motion* (PLM) about the economic dynamics, which takes an analogous form to the MSV solutions in (16). However, since the agents do not know the parameter values in matrices \bar{a} , \bar{b} and \bar{c} in (16), they rely on past data to make estimates a_t , b_t , and c_t each period, using a recursive least squares method described below. That is, at time t , agents' PLM, or what they perceive the economic dynamics to be is:

$$y_t = a_t + b_t y_{t-1} + c_t w_t \quad (17).$$

As our first case of knowledge imperfection, we assume the exogenous shock vector, w_t is observed by both the agents and the policymakers. In other words, compared to the information set under RE, $I_t^{RE} = \{\bar{a}, \bar{b}, \bar{c}, y^{t-1}, w^t\}$ where $z^t \equiv \{z_0, z_1, z_2, \dots, z_t\}$, the information set under this first case of learning is: $I_t^{AL1} = \{a_t, b_t, c_t, y^{t-1}, w^t\}$. At each time t , agents first rely on observed data up to and including period $t-1$ to update their parameter estimates a_t , b_t , and c_t . They also observe the value of the contemporaneous shocks, and form expectations using these information:

$$E_t y_{t+1} = a_t + b_t E_t y_t + c_t E_t w_{t+1},$$

and from (15) and (16):

$$E_t y_{t+1} = (I + b_t) a_t + b_t^2 y_{t-1} + (b_t c_t + c_t \rho_w) w_t. \quad (18)$$

where ρ_w is also assumed to be known by agents under this case of learning. At time t , the policymakers set the interest rate r_t according to their desired rules. The *Actual Law of Motion* (ALM) for y_t is generated according to (14) and (15), as the following:¹⁸

$$y_t = A + B[(I + b_t)a_t + b_t^2 y_{t-1} + (b_t c_t + c_t \rho_w)w_t] + C y_{t-1} + D w_t, \text{ or}$$

$$y_t = [A + B(I + b_t)a_t] + (B b_t^2 + C)y_{t-1} + [B(b_t c_t + c_t \rho_w) + D]w_t \quad (19).$$

At the beginning of next period $t + 1$, agents use the newly available information, e.g., data up to and including period t , to re-estimate the PLM and obtain new parameter estimates a_{t+1} , b_{t+1} and c_{t+1} . Once the shocks w_{t+1} are realized and observed, and the interest rate r_{t+1} is set by the policymakers, the ALM for y_{t+1} is generated and the learning process continues in this rolling fashion.

For estimating the parameters, we assume the adaptive learning agents to follow the recursive least squares algorithm given by:

$$\phi_t = \phi_{t-1} + g_t R_t^{-1} z_{t-1} (y_{t-1} - \phi_{t-1}' z_{t-1})' \quad (20)$$

$$R_t = R_{t-1} + g_t (z_{t-1} z_{t-1}' - R_{t-1})' \quad (21)$$

where $\phi_t = [a_t, b_t, c_t]'$, $z_t = [1, y_{t-1}, w_t]'$ and R_t is the updated matrix of second moments of the regressor z_t . The updating rate or the gain parameter, g_t , is a key parameter for the perpetual learning mechanism. This paper focuses on constant gain learning, where the gain parameter is a small constant positive number, $0 < g_t < 1$. The empirical macroeconomic literature suggests that the gain value for quarterly data is in the range of 0.01 and 0.04 (Milani (2007) and Chakraborty and Evans (2008)). Specifically, as discussed in Orphanides and Williams (2005), a

¹⁸ The ALM can be considered as the true data generating process. Also, the ALM is sometimes called the temporary equilibrium for endogenous variables. See Evans and Honkapojah (2006).

constant gain g indicates that agents use $2/g$ lags of data to form their expectations.¹⁹ We consider the gain value of 0.02, which implies that agents look at 25 years of historical data.

In sum, under this form of adaptive learning, the dynamics of the model are defined by the recursive least squares updating equations (20) and (21), the expectations formation process (18) derived from the PLM, the structural model equation (14), and the AR(1) process for the stochastic shocks w_t equation (15).

We next consider another setting where we introduce an additional imperfect knowledge distortion by assuming that private agents do not observe the contemporaneous structural shocks to the economy, w_t . With their information set further reduced to $I_t^{AL2} = \{a_t, b_t, c_t, y^{t-1}, w^{t-1}, \hat{w}_t\}$, they have to estimate the shocks at time t as well. Following Orphanides and Williams (2007a), we assume the agents use the following simple updating rule for the estimates of the three exogenous shocks in \hat{w}_t :

$$\begin{aligned}\hat{v}_t &= \hat{v}_{t-1} + 0.005(v_{t-1} - \hat{v}_{t-1}) \\ \hat{u}_t &= \hat{u}_{t-1} + 0.005(u_{t-1} - \hat{u}_{t-1}) \\ \hat{\pi}_t^* &= \hat{\pi}_{t-1}^* + 0.005(\pi_{t-1}^* - \hat{\pi}_{t-1}^*)\end{aligned}$$

where variables with ‘^’ represent the estimates of corresponding shocks. In addition to estimating a_t , b_t , and c_t using the recursive least square algorithm, in this second form of imperfect knowledge setting, agents also have to learn about the contemporaneous shocks to the economy.

By comparing the outcomes under the two scenarios of learning, we aim to assess the welfare benefit of knowing the current shocks, and in situations where it is not possible, whether monetary policy may need to adjust.

¹⁹ A smaller gain thus means that agents use longer history of past data to form their forecasts.

4. Numerical Analyses and Discussion

4.1 Calibration

For calibrating our benchmark models, we adopt parameter values from the following sources. Following GM (2005), we set $\sigma = \gamma = \eta = 1$. φ is set to be 3, which implies the elasticity of the supply of labor is 1/3. We set the degree of openness parameter α to be 0.4, corresponding to the share of imports to GDP in Canada (a small open economy). Parameter β is set to be 0.99, which implies that a riskless annual return (ρ) is 4 percent. We set $\theta = 0.75$, corresponding to an average period of four quarters between price adjustments. As for the parameters governing monetary policy rules, we set $\varphi_\pi = \varphi_{\pi_H} = \varphi_{\pi^*} = 1.5$ and $\varphi_x = 0.5$ as in Taylor (1993). As a robustness check, we explore various alternative policy weights, $\varphi_\pi, \varphi_{\pi_H} \in \{1, 1.5, 2\}$ and $\varphi_x \in \{0, 0.5, 0.75, 1\}$.²⁰ Overall we find the same qualitative conclusions and, overall, the standard Taylor weights outperform these alternatives. The targets of DI (π_H^T) and CI (π^T) are set to be zero.

We assume the three stochastic shocks $\{v_t, u_t, \pi_t^*\}$ follow independent AR (1) processes, and consider three sets of values: i) Baseline; ii) Less Volatile Foreign (inflation) Shock; and iii) Less Volatile Domestic (cost-push) Shock.²¹ For the productivity shock, which we keep the same across the three scenarios, we first fit an AR (1) process to the (log) labor productivity index of Canada, using OECD Economic Outlook data over the sample period of 1976:1-2006:2. We obtain $v_t = 0.82v_{t-1} + \varepsilon_{v,t}$ with $\sigma_{\varepsilon_{v,t}} = 0.005$. We then examine the case where the standard deviation is twice as large, $\sigma_{\varepsilon_{v,t}} = 0.01$, and upon confirming the same qualitative results, use it as our productivity shock process. While the key qualitative conclusions are the same between these two choices, the more volatile shock, when used in combination with the other two scenarios with differing inflation volatilities, illustrate our points more clearly. We believe this is a reasonable adjustment as Canada is one of the more stable small open economies in the

²⁰ We report selective results in Table 5.

²¹ We name these scenarios this way because in cases ii) and iii), we reduce the volatility of the foreign or the domestic nominal shock 50-fold respectively, from 5.00E-3 to 1.00E-4.

world. For the domestic cost push shock, we assume the following process: $u_t = 0.4u_{t-1} + \varepsilon_{u,t}$ with $\sigma_{\varepsilon_{u,t}} = 0.005$, as in Evans and Honkapohja (2003). Lastly, for the world inflation process π_t^* , we fit an AR(1) processes to (log) U.S. CPI over the same sample period, and obtain $\pi_t^* = 0.92\pi_{t-1}^* + \varepsilon_{\pi^*,t}$, with $\sigma_{\varepsilon_{\pi^*,t}} = 0.005$. This is our baseline setup. To explore how monetary policy should adjust in response to inflationary shocks of different volatility originating from home versus abroad, we consider the following sets of parameters. For the case of ‘Less Volatile Foreign Shocks,’ we set the standard deviation of the white noise terms in the foreign shock, $\sigma_{\varepsilon_{\pi^*,t}}$ to be 0.0001, keeping everything else the same as the benchmark case. For the less ‘Less Volatile Home Shocks’ case, we set $\sigma_{\varepsilon_{u,t}} = 0.0001$ instead.

4.2 Simulation Results

We conduct simulation experiments to compare welfare losses associated with two monetary policy rules: F-DITR and F-CITR. The learning algorithm is based on Evans and Honkapohja (2001, 2006). We also follow Orphanides and Williams (2007a) in imposing a “projection facility” to keep the simulation paths non-explosive, and we provide more descriptions in Appendix.²² We simulate the dynamics of the economy 300 times for 450 periods each and discard the first 50 periods to reduce the effects of initial conditions.

Table 1 provides a baseline calibration for the welfare losses associated with F-DITR and F-CITR under RE and under the two learning setups where shocks $\{v_t, u_t, \pi_t^*\}$ are observable and unobservable. Comparing across the three sets of two-columns, we first note that when knowledge imperfection in the form of ignorance over the parameter values governing the dynamic equations does not incur much additional welfare cost. That is, learning with observable shocks produce welfare results essentially the same as those obtained under RE. This coincides with findings in Williams (2003), in the closed economy learning literature. Under the baseline shocks assumed here, we also see that targeting domestic inflation (F-DITR) leads to lower welfare loss in general. We next consider two alternative structural parameters to confirm

²² The “projection facility” has often been imposed in the learning literature. See, for example, Chakraborty and Evans (2008) and Waters (2009).

that the behavior of our model makes economic sense. First, when firms adjust prices less often, i.e., when θ , the probability of having to keep the old price each period increases from the benchmark 0.75 to 0.9, the distortions in prices and in output allocation are more severe. Correspondingly, we see welfare losses being amplified by approximately eight-fold or more from the benchmark case. Next, when a country is less open to trade (import share α drops from 0.4 or 0.1), the welfare impacts of F-DITR and F-CITR become very similar, since CPI and the domestic price level converge.

In Table 2, we vary the volatility of the shocks to see how policy conclusions compare under RE and learning in response to the relative size of shocks. As described earlier, the two additional cases are: 1) ‘Less Volatile Foreign Shocks’, with the volatility of π_t^* reduced, and 2) ‘Less Volatile Home Shocks’, where the variance of the domestic nominal shocks (u_t) is reduced. We compare the RE outcome with the learning results where agents do not observe contemporaneous shocks (the middle two columns), and also where they do have this information (last two columns). From the first two columns, we first see that under RE, the relative size of the shocks can alter the preferred policy choice (denoted with an underline.) When foreign shocks are much less volatile, F-CITR is preferable. By targeting CI instead of DI, the policymakers are bringing foreign shocks into the domestic macroeconomic dynamics. As discussed earlier, quieter foreign shocks lead to more stable terms of trade, so the expenditure-switching channel has only a small perverse effect on the volatility of domestic output (Table 3 shows that output volatility increased from 0.3074 to 0.3411.) However, with the stable foreign inflation rate as a nominal anchor, domestic inflation variation is substantially reduced, from 0.8933 to 0.7381, as shown in Table 3 also. Since the welfare loss function penalizes inflation volatility more heavily relative to output volatility, F-CITR policy rule delivers higher welfare. That is, the inflation anchor effect outweighs the small expenditure switching effect. We can view this result in under signal extraction framework, as discussed in Mankiw and Reis (2003) who argue that inflation-targeting central banks should assign low target weights to the food and energy sectors which have highly volatile sector-specific shocks. With stable foreign shocks, the noise in the terms of trade is reduced. Targeting the terms of trade dynamics therefore helps extract the relevant signal.

We next evaluate the performances of the two policy rules when the domestic cost-push shocks are less volatile, relative to foreign shocks (the ‘Less Volatile Home Shocks’ case). F-DITR, which puts more weight on stabilizing DI, significantly dampens the losses, from 0.1680 with baseline shocks to 0.0388 with less volatile home shocks. In sum, we see that under RE, policymakers should monitor variables that experience less volatile shocks.

The middle two columns in Table 2 report the simulation results under the scenario where adaptive learning agents *can* observe contemporaneous shocks, $\{v_t, u_t, \pi_t^*\}$. We note that the welfare outcomes do not differ much from that under RE. Targeting away from the more volatile shocks helps anchor expectations and improve welfare, as in the RE framework.

The last two columns in Table 2 show the welfare results under the case of learning when agents have limited knowledge and have to learn the dynamics governing *both* the relevant economic variables *and* the underlying structural shocks. We first note that the presence of the additional source of imperfect knowledge induces significant distortion in the economy. Unlike the previous milder form of knowledge imperfection discussed above, the welfare loss is much higher in all cases compared to that under the RE. Moreover, in contrast to the previous policy conclusions, here F-DITR is the better policy rule regardless of the source of volatile shocks. As is evident from Table 3, under RE, the F-DITR strategy, which leaves the economy isolated from the foreign shocks, is the preferred rule except when the foreign shock can serve as an anchor, i.e. when it is comparatively stable. When the shocks are unknown and agents have to learn about them, this anchoring advantage is no longer evident, as it becomes dominated by the uncertainty and forecast errors generated by having to learn about an additional variable. As shown in the last two rows of Table 3, this extra volatility induced by learning is especially relevant when the foreign shock is volatile. Under this more severe form of knowledge imperfection, F-DITR, which does not (implicitly or explicitly) respond to the unobservable foreign shocks leads to less volatility output and DI compared to F-CITR and helps stabilize the economy better.

The additional source of imperfect knowledge alters the RE policy conclusion. Even when foreign shocks are very quiet, the preferred policy rule is still to target DI (F-DITR). Put it differently, in situations where agents lack information about current shocks in addition to the

structural dynamics governing the economy, monetary policy *can* improve welfare by deviating from the RE policy prescription.

The different welfare and policy conclusions we obtain for the two different types of common adaptive learning scenarios are directly linked to the forecast errors generated in agents' learning process, which feeds into aggregate volatility and therefore welfare. Table 4 looks at these forecast errors more carefully. We compute the root mean square deviations (RMSD) to measure the deviations of output gap and inflation variables under learning relative to the ones under RE, as follows:

$$\text{RMSD} = \sqrt{\sum_{t=1}^N (X_t^{\text{Learning}} - X_t^{\text{RE}})^2 / N},$$

where, at time t , X_t^{Learning} is the value of a variable under learning and X_t^{RE} is the value of the same variable under RE. N is the number of simulations, which we set to 300. Table 4 reports the average of RMSDs over the 300 simulations under the two types of learning. Comparing Panel A and B, we see that when agents lack information about current structural shocks, RMSD are a thousand-fold higher than when they can observe these shocks. Imperfect knowledge about the current shocks pushes the economy substantially away from RE. This observation tells us that policymakers can correct the distortion from imperfect knowledge and improve welfare substantially, if they can provide information about contemporaneous structural shocks. The imperfect knowledge distortions associated with the structural parameters, on the other hand, are much less serious. As shown earlier, they do not alter the policy prescription either (between targeting DI versus CI).

As a robustness check to our conclusions based on the Taylor rules, we also compare policy performance of a pure inflation targeting rule.²³ We consider cases when the policymaker only targets either expected DI or CI, without the reacting to the output gap. Table 5 reports the welfare losses associated with the four forward-looking rules: a pure F-DI targeting rule, F-DITR,

²³ We also explored alternative weights on the output gap and inflation term in the context of the Taylor rule. Our qualitative conclusions are robust to these variations. All these rules have associated welfare losses and they are not optimal. As explained earlier, we choose these simple rules as they are more practical and may reflect actual policymaking more accurately.

a pure F-CI targeting rule, and F-CITR. Comparing between F-DI and F-CI, we observe the same conclusions as under the Taylor rules discussed earlier. First, when agents can observe contemporaneous shocks, distortions created from the learning process are not severe, as can be seen from the first two lines of results in each of the three panels. In addition, we see that under the mild form of imperfect knowledge, the policy comparison between targeting DI (F-DI) versus CI (F-CI) result in the same conclusions as under RE. But when shocks are not observable, the policymaker should switch to targeting DI, even when foreign shocks are relatively stable. As in the previous Taylor Rule analysis, this is contrary to the policy conclusion under RE.

It is also worth noting that when comparing the results between pure inflation targeting and the forward-looking Taylor rules, the Taylor rule is preferred in most cases. Especially in the domestic inflation-based policy rule scenarios in the first two columns, the Taylor rule results dominate in all cases, regardless of the relative variances of the shocks or RE or learning. Since the welfare function depends on the variances of both the output gap and DI, some degree of output targeting works better than a pure inflation target. In the CPI-based scenarios where the policymakers bring in foreign shocks into the interactions, the results are not as clear-cut. Overall, we find that when agents cannot observe contemporaneous structural shocks, explicitly targeting the output gap in addition to DI (CI) improve welfare, that is, F-DITR (F-CITR) dominates the purely F-DI (F-CI) targeting rule.

5. Conclusion

This paper incorporates adaptive learning into a standard New-Keynesian open economy dynamic stochastic general equilibrium (DSGE) model and analyze under what conditions policymakers should target *domestic producer price* inflation (DI) versus *consumer price* inflation (CI). Our goal is to examine how monetary policy rules should adjust when agents' information sets deviate from those assumed under the rational expectation paradigm. We assume the policymaker follows a forward-looking Taylor rule and focus on analyzing the interplay between the source of the dominant shocks and the extent of knowledge imperfection.

We find that even though the central bank has no informational advantage, monetary policy can nonetheless facilitate the learning process and mitigate distortions associated with imperfect knowledge. Specifically, when agents have very limited knowledge and have to learn the dynamics governing *both* the relevant economic indicators *and* the underlying structural shocks, a DI target-based Taylor rule produces smaller forecast errors and is thus better at stabilizing the economy. However, when agents can observe contemporaneous shocks and need only to learn how key economic variables evolve (a situation akin to a post-structural-shift economy), targeting away from the dominant shocks helps anchor expectations and improve welfare, as in the rational expectation framework. Lastly, we find that the relatively cost of knowledge imperfection, in terms of the excess volatility induced in the economy, is much more severe when agents lack information about current shocks, pointing to the importance of information dissemination and transparent policy-making.

References

- Aoki, K., 2001. Optimal Monetary Policy Responses to Relative-Price Changes. *Journal of Monetary Economics* 48, 55--80.
- Benigno, G., Benigno P., 2003. Price Stability in Open Economies. *Review of Economic Studies* 70, 743--764.
- Bullard, J., Schaling, E., 2006. Monetary Policy, Determinacy, and Learnability in a Two-Block World Economy. Federal Reserve Bank of St. Louis Working Paper 2006-038A.
- Calvo, G., 1983. Staggered Prices in a Utility Maximizing Framework. *Journal of Monetary Economics* 12, 383--398.
- Campolmi, A., 2008. Which Inflation to Target? A Small Open Economy with Sticky Wages. Working Paper, Central European University.
- Chakraborty, A., Evans, G., 2008. Can Perpetual Learning Explain the Forward Premium Puzzle? *Journal of Monetary Economics* 55, 477--490.
- Cogley, T., Primiceri, G., and Sargent, T.J., 2008. Inflation-Gap Persistence in the U.S.. NBER Working Paper 13749.
- Clarida, R., Gali, J., and Gertler, M., 1999. The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature* 37, 1661--1707.
- Clarida, R., Gali, J., and Gertler, M., 2001. Optimal Monetary Policy in Open versus Closed Economies: An Integrated Approach. *The American Economic Review* 91, 148--252.
- Clarida, R., Gali, J., and Gertler, M., 2002. A Simple Framework for International Monetary Policy Analysis. *Journal of Monetary Economics* 49, 879--904.
- Corsetti, G., Pesenti, P., 2001. Welfare and Macroeconomic Interdependence. *Quarterly Journal of Economics* 116, 421--446.
- Corsetti, G., Pesenti, P., 2005. International Dimensions of Optimal Monetary Policy. *Journal of Monetary Economics* 52, 281--305.
- Doyle, B. M., Erceg, C. J., and Levin, A. T., 2006. Monetary Policy Rules in Economies with Traded and Non-traded Goods. Unpublished manuscript, Federal Reserve Board.

- Evans, G., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.
- Evans, G., Honkapohja, S., 2003. Adaptive Learning and Monetary Policy Design. *Journal of Money, Credit and Banking* 35, 1045--1072.
- Evans, G., Honkapohja, S., 2006. Monetary Policy, Expectations and Commitment. *Scandinavian Journal of Economics* 108, 15--38.
- Gali, J., 2003. New Perspectives on Monetary Policy, Inflation, and the Business Cycle, in: Dewatripont, M., Hansen, L., and Turnovsky, S. (Eds.), *Advances in Economics and Econometrics*, Vol. 3. Cambridge University Press, Cambridge, pp.151-197.
- Gali, J., Monacelli, T., 2005. Monetary Policy and Exchange Rate Volatility in A Small Open Economy. *Review of Economic Studies* 72, 707--734.
- Kim, C-J., Nelson, C., 1999. Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of Business Cycle. *Review of Economics and Statistics* 81, 608--616.
- Llosa, G., Tuesta, V., 2008. Determinacy and Learnability of Monetary Policy Rules in Small Open Economies. *Journal of Monetary Economics* 40, 1033--1063.
- Mankiw, N. G., Reis, R., 2003. What measure of Inflation Should a Central Bank Target? *Journal of the European Economic Association* 1, 1058--1086.
- Mark, N. C., 2009. Changing Monetary Policy Rules, Learning, and Real Exchange Dynamics, *Journal of Money, Credit and Banking*, forthcoming.
- McCallum, B. T., 1983. On Non-Uniqueness in Rational Expectations Models: An Attempt at Perspective. *Journal of Monetary Economics* 11, 134--168.
- McCallum, B. T., 1998. Solutions to Linear Rational Expectations Models: A Compact Exposition. *Economics Letters* 61, 143--147.
- Milani, F., 2007. Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics* 54, 2065--2082.

- Orphanides, A., Williams, J. C., 2005. Imperfect Knowledge, Inflation Expectations and Monetary Policy, in: Bernanke, M., Woodford, M. (Eds.), *The Inflation Targeting Debate*, University of Chicago Press, Chicago, IL, pp. 201-246.
- Orphanides, A., Williams, J. C., 2007a. Inflation Targeting under Imperfect Knowledge, in: Mishkin, F., Schmidt-Hebbel, K. (Eds.), *Monetary Policy under Inflation Targeting*, Central Bank of Chile, Santiago, Chile, pp.77-123.
- Orphanides, A., Williams, J. C., 2007b. Robust Monetary Policy with Imperfect Knowledge. *Journal of Monetary Economics* 54, 1406--1435.
- Preston, B., 2005. Learning about Monetary Policy Rules when Long-Horizon Expectations Matter. *International Journal of Central Banking* 1, 81--126.
- Preston, B., 2008. Adaptive Learning and the Use of Forecasts in Monetary Policy. *Journal of Economic and Dynamics Control* 32, 3661--3681.
- Sargent, T. J., 1993. *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford, NY.
- Stocks, J., Watson, M., 2005. Understanding Changes in International Business Cycle Dynamics. *Journal of the European Economic Association* 3, 968--1006.
- Svensson, L.E.O., 2000. Open-economy Inflation Targeting. *Journal of International Economics* 50, 155--183.
- Taylor, J. B., 1993. Discretion versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy* 39, 195--214.
- Waters, G., 2009. Learning, Commitment and Monetary Policy: the Case for Partial Commitment. *Macroeconomic Dynamics*, forthcoming.
- Williams, N., 2003. Adaptive Learning and Business Cycles, Mimeo. University of Wisconsin-Madison.
- Woodford, M., 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.

Appendix: The Implementation of the Learning Algorithms

As initial conditions for each adaptive learning simulation, we perturb the rational expectations equilibrium with a small white noise as follows: $a = \bar{a} + 0.005 \times \text{random}$, $b = \bar{b} + 0.004 \times \text{random}$, $c = \bar{c} + 0.002 \times \text{random}$, where “random” is an innovation generated from a uniform distribution. We set $y_0 = \bar{y}$ and $R_0 = \bar{R}$. The innovations in each period are drawn from normal distributions.

In addition, to keep the stochastic simulation non-explosive, we implement two additional algorithms suggested by Orphanides and Willaims (2007a) to reflect the view that in practice, private agents would reject unstable models so our analyses should similarly rule them out. In each period, we compute the roots of the modulus of the forecasting VAR, excluding the constants. If all of the roots are in the modulus of 1, the forecast model is updated as discussed in the text. If not, the forecast model is not updated and the matrices ϕ and R are kept at their respective values from the previous period. We further impose the following condition to restrain explosive behavior: if any of the relevant variables exceeds, in absolute value, five times its unconditional standard deviations (computed under the assumption of rational expectations), then the variable that exceeds this bound is set to the boundary value for that period. However, these two constraints are not sufficient to rule out all explosive behavior in our adaptive learning simulations. Thus, we compute relevant statistics using only simulation runs that give variables standard deviations that are less than ten times their respective unconditional standard deviations under rational expectations.

Table 1: Welfare Losses under F-DITR and F-CITR with Baseline Calibration for Shocks for RE and Learning with Unobservable Shocks and Observable Shocks

	RE		Learning with Observable Shocks		Learning with Unobservable Shocks	
	F-DITR	F-CITR	F-DITR	F-CITR	F-DITR	F-CITR
Benchmark	<u>0.1680</u>	0.1893	<u>0.1682</u>	0.1884	<u>0.4182</u>	4.5889
Parameter	(0.0010)	(0.0016)	(0.0010)	(0.0013)	(0.0393)	(0.3155)
Price Stickier ($\theta=0.90$)	<u>1.1714</u>	1.2209	<u>1.1749</u>	1.2129	<u>4.5117</u>	43.0907
	(0.0061)	(0.0074)	(0.0059)	(0.0063)	(0.4355)	(2.3347)
Less Openness ($\alpha=0.10$)	0.2547	<u>0.2416</u>	0.2546	<u>0.2421</u>	<u>0.6954</u>	1.6082
	(0.0016)	(0.0017)	(0.0015)	(0.0014)	(0.0678)	(0.1323)

Note: Numbers reported are the averaged expected welfare loss, multiplied by 100, over 300 simulations. The welfare loss is measured as annualized percentage deviation from optimal steady state consumption. The numbers in parentheses are the standard errors of statistics. We omit results that do not satisfy the projection facility conditions, as discussed in the Appendix. Underlined numbers represent the policy choice with lower welfare loss within each scenario; we note that they are not always statistically lower.

Table 2: Welfare Losses under F-DITR and F-CITR for RE and Learning with observable Shocks and Unobservable Shocks

	RE		Learning with Observable Shocks		Learning with Unobservable Shocks	
	F-DITR	F-CITR	F-DITR	F-CITR	F-DITR	F-CITR
Baseline Shocks	<u>0.1680</u> (0.0010)	0.1893 (0.0016)	<u>0.1682</u> (0.0010)	0.1884 (0.0013)	<u>0.4182</u> (0.0393)	4.5889 (0.3155)
Less Volatile Foreign Shocks	0.1689 (0.0011)	<u>0.1159</u> (0.0008)	0.2507 (0.0211)	<u>0.1148</u> (0.0006)	<u>0.7381</u> (0.0665)	1.1729 (0.0732)
Less Volatile Home Shocks	<u>0.0388</u> (0.0003)	0.0834 (0.0014)	<u>0.0425</u> (0.0023)	0.0854 (0.0052)	<u>0.1222</u> (0.0119)	4.9196 (0.3049)

Note: Numbers reported are the averaged expected welfare loss, multiplied by 100, over 300 simulations. The welfare loss is measured as annualized percentage deviation from optimal steady state consumption. The numbers in parentheses are the standard errors of statistics. We omit results that do not satisfy the projection facility conditions, as discussed in the Appendix. Underlined numbers represent the preferred policy choice within each scenario.

Table 3: Standard Deviations of the Output Gap and Domestic Inflation under F-DITR and F-CITR for RE and Learning with Unobservable Shocks

	RE				Learning with Unobservable Shocks			
	F-DITR		F-CITR		F-DITR		F-CITR	
	x_t	$\pi_{H,t}$	x_t	$\pi_{H,t}$	x_t	$\pi_{H,t}$	x_t	$\pi_{H,t}$
Baseline Shocks	0.3054 (0.0011)	0.8910 (0.0027)	0.3719 (0.0013)	0.9442 (0.0040)	0.9431 (0.0262)	1.1926 (0.0432)	1.3049 (0.0314)	4.0429 (0.1639)
Less Volatile Foreign Shocks	0.3074 (0.0012)	0.8933 (0.0030)	0.3411 (0.0010)	0.7381 (0.0024)	1.0825 (0.0263)	1.5437 (0.0694)	1.2613 (0.0317)	2.0607 (0.0803)
Less Volatile Home Shocks	0.2331 (0.0010)	0.4254 (0.0019)	0.1847 (0.0013)	0.6196 (0.0047)	0.4514 (0.0176)	0.6193 (0.0251)	0.6656 (0.0152)	4.3034 (0.1616)

Note: Numbers reported are the averaged statistics over 300 simulations. The parenthesized numbers are the standard errors of statistics.

Table 4:

A) Root Mean Square Deviations of Key Variables under F-DITR and F-CITR for Learning with Observable Shocks

	x_t		$\pi_{H,t}$		π_t	
	F-DITR	F-CITR	F-DITR	F-CITR	F-DITR	F-CITR
Baseline Shocks	5.39 e-6	3.71 e-6	1.57 e-6	2.63 e-6	1.61 e-6	2.73 e-6
Less Volatile Foreign Shocks	2.25 e-5	3.79 e-6	1.28 e-5	2.68 e-6	1.16 e-5	2.72 e-6
Less Volatile Home Shocks	9.26 e-6	5.98 e-6	5.39 e-6	4.00 e-6	5.65 e-5	7.14 e-6

B) Root Mean Square Deviations of Key Variables under F-DITR and F-CITR for Learning with Unobservable Shocks

	x_t		$\pi_{H,t}$		π_t	
	F-DITR	F-CITR	F-DITR	F-CITR	F-DITR	F-CITR
Baseline Shocks	3.02 e-3	8.44 e-3	7.32 e-3	3.31 e-3	2.16 e-3	7.06 e-3
Less Volatile Foreign Shocks	1.26 e-2	1.95 e-2	1.81 e-2	1.78 e-3	2.18 e-3	1.27 e-2
Less Volatile Home Shocks	1.44 e-3	9.88 e-3	5.42 e-3	3.11 e-3	1.05 e-3	9.67 e-3

Note: Numbers reported are the averaged RMSD over 300 simulations. We omit results that do not satisfy the projection facility conditions, as discussed in the Appendix.

Table 5: Welfare Losses under Purely F-DI Targeting, F-DITR, Purely F-CI Targeting, and F-CITR for RE and Learning with Unobservable Shocks and Observable Shocks

	Purely F-DI Targeting	F-DITR	Purely F-CI Targeting	F-CITR
1) Baseline Shocks				
RE	0.1997 (0.0014)	0.1680 (0.0010)	0.2131 (0.0023)	0.1893 (0.0016)
Learning with Observable Shocks	0.2030 (0.0014)	0.1682 (0.0010)	0.2142 (0.0018)	0.1884 (0.0013)
Learning with Unobservable Shocks	0.6024 (0.0603)	0.4182 (0.0393)	4.7603 (0.3263)	4.5889 (0.3115)
2) Less Volatile Foreign Shocks				
RE	0.2016 (0.0014)	0.1689 (0.0011)	0.0976 (0.0007)	0.1159 (0.0008)
Learning with Observable Shocks	0.2795 (0.0233)	0.2507 (0.0211)	0.0976 (0.0006)	0.1148 (0.0006)
Learning with Unobservable Shocks	0.9668 (0.0785)	0.7381 (0.0665)	1.2600 (0.0643)	1.1729 (0.0732)
3) Less Volatile Domestic Shocks				
RE	0.0810 (0.0008)	0.0388 (0.0003)	0.1352 (0.0021)	0.0834 (0.0014)
Learning with Observable Shocks	0.0860 (0.0063)	0.0425 (0.0023)	0.1336 (0.0017)	0.0854 (0.0052)
Learning with Unobservable Shocks	0.3252 (0.0309)	0.1222 (0.0119)	5.4559 (0.3222)	4.9196 (0.3049)

Note: Numbers reported are the averaged expected welfare loss, multiplied by 100, over 300 simulations. The welfare loss is measured as annualized percentage deviation from optimal steady state consumption. The numbers in parentheses are the standard errors of statistics. We omit results that do not satisfy the projection facility conditions, as discussed in the Appendix.