Interest Rate Rules in DSGE Models: Tracking the Efficient Real Interest Rate☆

Vasco Cúrdia[a], Andrea Ferrero[b], Ging Cee Ng[c], Andrea Tambalotti[b,*]

March 2012

[a]*Federal Reserve Bank of San Francisco*
[b]*Federal Reserve Bank of New York*
[c]*University of Chicago*

Abstract

The empirical Dynamic Stochastic General Equilibrium (DSGE) literature pays surprisingly little attention to the behavior of the monetary authority. Alternative policy rule specifications abound, but their relative merit is rarely discussed. We contribute to filling this gap by undertaking a systematic comparison of the fit of several families of interest rate rules within two popular New Keynesian models. We find that the best fitting rules are those in which the nominal interest rate tracks the evolution of the model-consistent efficient real interest rate—an empirical hypothesis not previously explored in the literature.

Keywords: Taylor rules, Efficient real interest rate, Bayesian model comparison

JEL Classification: E43, E58, C11

1. Introduction

Most central banks around the World pursue some form of flexible inflation targeting, striving to stabilize inflation while keeping a watchful eye on real economic developments. In the United States, the Federal Reserve Act requires the monetary authority to pursue the dual mandate of stabilizing prices and promoting "maximum employment." In a seminal contribution, Taylor (1993) proposed a parsimonious approach to capturing the implications

☆We thank Alejandro Justiniano, Robert King, Thomas Lubik, Giorgio Primiceri, Keith Sill, Chris Sims, an anonymous referee and many seminar and conference participants for comments and suggestions. Matthew Hom provided excellent research assistance. The views expressed in this paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of San Francisco, the Federal Reserve Bank of New York or the Federal Reserve System.

*Corresponding Author: Federal Reserve Bank of New York, 33 Liberty Street, 3rd Floor, New York, NY 10045. Phone: (212) 720-5657; Fax: (212) 720-1844

Email addresses: `vasco.curdia@sf.frb.org` (Vasco Cúrdia), `andrea.ferrero@ny.frb.org` (Andrea Ferrero), `gingceeng@gmail.com` (Ging Cee Ng), `andrea.tambalotti@ny.frb.org` (Andrea Tambalotti)

of this dual objective for central banks' interest rate decisions—the now famous Taylor rule. According to this policy rule, the monetary authority sets the nominal interest rate in response to deviations of inflation from its target and to some measure of real economic activity, such as the output gap. Although originally proposed as a simple normative guide for monetary policy, the observation that the rule tracked the actual behavior of the Fed under Chairman Greenspan between 1987 and 1992 quite closely turned interest rate feedback rules into an ubiquitous positive tool used to describe actual policy decisions. This evolution of Taylor rules from normative to positive devices has been supported by a large literature documenting their excellent empirical performance using single equation estimation methods (Judd and Rudebusch, 1998; Clarida et al., 1999, 2000; English et al., 2003; Orphanides, 2003 and, most recently, Coibion and Gorodnichenko, 2011).

Their parsimony, convenience, and empirical fit has also made Taylor-type rules the most common formulation for the description of monetary policy within dynamic stochastic general equilibrium (DSGE) models. Most DSGE studies, however, simply posit an interest rate feedback rule broadly inspired by Taylor (1993), usually with no discussion of its details and potential alternatives. This casual treatment of the behavior of the monetary authority in DSGE papers stands in stark contrast with the attention more generally dedicated to the modeling of the private sector. As a result, this literature has witnessed a proliferation of policy specifications—especially with respect to the choice of the real variables the central bank reacts to—but with very little guidance on their empirical performance within the model under consideration.[1]

The research summarized in this paper represents an attempt to put some order in this chaotic situation by systematically comparing the fit of a large set of interest rate rules

[1]See Schmitt-Grohé and Uribe (2007) for a recent normative analysis of alternative simple interest rate rules within a calibrated DSGE model. Svensson (2003) recommends modeling central banks as optimizing agents that maximize an objective function, as customary for the private sector, rather than as automatons committed to an interest rate feedback rule. The optimal targeting rule obtained in this framework, however, still depends on the arguments of the loss function policymakers are assumed to minimize. See Adolfson et al. (2008) for a state-of-the-art implementation of this approach within a DSGE model for Sweden.

within an estimated DSGE model of the U.S. economy. Most of the rules we analyzed have previously appeared in the literature. Others, including the best-fitting one, have not. The most notable feature of the best-fitting rule is that the monetary authority tracks the evolution of the efficient real interest rate—the real interest rate that would prevail in equilibrium if the economy were perfectly competitive—as an indicator of real economic developments. In fact, this measure of the equilibrium interest rate is a better proxy for the real economic developments to which monetary policy seems to respond than any of the several measures of output gap we experimented with. This result sets our contribution apart from the literature on single-equation estimation of Taylor rules, as we need a complete general equilibrium model to compute the measure of real interest rate analyzed here.

This policy rule echos Wicksell's suggestion that a "natural" rate of return determined by real factors represents a useful target for monetary policy (Woodford 2003), an idea familiar to Fed policymakers at least since the early 1990s (e.g. Greenspan, 1993). To our knowledge, we are the first to estimate interest rate rules consistent with this idea and to document their strong empirical performance.[2]

Our extensive investigation of the empirical properties of these rules within a DSGE framework suggests that the inclusion of the efficient real interest rate as a time-varying intercept in the policy equation improves model fit across different specifications of the policy rule and of the private sector behavior. In its current version, the paper documents this fact by focusing on a few especially representative cases. Cúrdia et al. (2011b) show that similar results hold across a much wider spectrum of policy specifications, which includes most of the interest rate equations posited in the DSGE literature and those estimated with single equation methods.

Methodologically, we follow a Bayesian empirical strategy. First, we embed each of the policy rules whose fit we wish to evaluate within a DSGE framework, defined by given tastes and technology describing the private sector. This step produces a set of models, one for

[2] Trehan and Wu (2007) discuss the biases in the reduced-form estimation of policy rules with a constant intercept, when in fact the central bank responds to a time-varying equilibrium real rate, but do not estimate this response.

each different policy specification. We then estimate each of these models with Bayesian methods and compare their fit using marginal data densities.[3] This criterion produces a ranking of the different models, as well as a quantitative measure of their relative ability to fit the data. Importantly, the resulting measure evaluates the congruence of the entire model with the data, rather than of the policy equation alone. In a general equilibrium context, this approach provides an fully coherent ranking, although we also look at other, more informal, indicators of the extent to which different policy rules make the estimated model more or less "reasonable."

Overall, our findings suggest that the specification of the interest rate rule can have a significant impact on the fit of DSGE models. The gap in marginal likelihoods between the best and worst fitting rules ranges between fifty and eighty log-points depending on the model under consideration. As a reference point, these differences in fit are of the same order of magnitude as those between similar structural models estimated with or without stochastic volatility (Cúrdia et al., 2011a). This evidence underscores the importance for DSGE researchers of paying significantly more attention to the specification of monetary policy than is common practice to date.

The rest of the paper proceeds as follows. Section 2 presents a small-scale model of private sector behavior together with the baseline interest rate rule. Section 3 discusses the econometric methodology and the estimation results for the baseline model. Section 4 introduces the alternative classes of policy rules we consider and compares their empirical performance. Section 5 extends the analysis to a medium-scale model. Section 6 concludes.

2. A Simple Model of the Monetary Transmission Mechanism

We augment the purely forward-looking textbook New Keynesian framework (Woodford, 2003) with two sources of inertia to improve its ability to fit the data. On the demand side,

[3] An and Schorfheide (2007) provide a comprehensive survey of the application of Bayesian methods to the estimation and comparison of DSGE models. Lubik and Schorfheide (2007) use similar methods to estimate the response of monetary policy to exchange rate movements in several small open economies.

we include habits in consumption in the utility specification. On the supply side, we allow for partial indexation to past inflation of the subset of prices that are not reoptimized in each period. We also allow for exogenous growth in total factor productivity and shocks to its growth rate.

The resulting model is smaller than the workhorse empirical DSGE model, such as Smets and Wouters (2007). In particular, we abstract from capital accumulation and the attending frictions (endogenous utilization and investment adjustment costs) and from non-competitive features in the labor market (monopolistic competition and sticky wages). This modeling choice allowed us to estimate and compare the fit of a very large number of interest rate rules at various stages of this project, without having to worry about computational constraints. In this paper, we only present the highlights of the results. Cúrdia et al. (2011b) report results for all the rules—55 in total—and their detailed ranking in terms of fit.

The remainder of this section presents the linearized equilibrium conditions of the model, which constitute the basis for estimation. Appendix A contains details of the model's microfoundations, including the mapping of the tastes and technology parameters into those of the approximate log-linear equations.

2.1. Private Sector

The model consists of aggregate demand and supply relations and a monetary policy rule, plus some auxiliary equations to define economic variables that will be useful later in the interpretation of the results. We present the model in "gap" form. The main benefit of this approach is that the resulting structure indeed takes the very familiar representation of the baseline New Keynesian model.

The aggregate demand side of the model is an Euler equation in the measure of real activity $\tilde{x}_t$

$$\tilde{x}_t = E_t\tilde{x}_{t+1} - \varphi_\gamma^{-1}\left(i_t - E_t\pi_{t+1} - r_t^e\right), \tag{1}$$

where i_t is the (continuously compounded) nominal interest rate, π_t is inflation, r_t^e is the efficient real interest rate, and $E_t(.)$ is the expectation operator conditional on all available information at time t. The parameter φ_γ measures the sensitivity of real activity to the real

interest rate (gap).

The real activity measure in expression (1) is a distributed lag of the efficient output gap $x_t^e \equiv y_t - y_t^e$

$$\tilde{x}_t \equiv \left(x_t^e - \eta_\gamma x_{t-1}^e\right) - \beta\eta_\gamma E_t\left(x_{t+1}^e - \eta_\gamma x_t^e\right),$$

where y_t is real output and y_t^e is its efficient counterpart. The lead-lag structure in the definition of $\tilde{x}_t$ reflects the presence of internal habits in consumption, to a degree indexed by the parameter η_γ.

The efficient level of output y_t^e and especially the efficient real interest rate r_t^e are key constructs in our analysis below. Efficient output represents the level of aggregate production that would prevail in equilibrium if prices were—and always had been—flexible, absent any markup shocks. Efficient output evolves according to the difference equation

$$\omega y_t^e + \varphi_\gamma\left(y_t^e - \eta_\gamma y_{t-1}^e\right) + \beta\varphi_\gamma\eta_\gamma\left(E_t y_{t+1}^e - \eta_\gamma y_t^e\right) = \varphi_\gamma\eta_\gamma\left(\beta E_t\gamma_{t+1} - \gamma_t\right) + \frac{\beta\eta_\gamma}{1-\beta\eta_\gamma}E_t\delta_{t+1}, \quad (2)$$

from which we observe that y_t^e is a linear combination of the past and future expected values of productivity growth γ_t and intertemporal taste shocks δ_t. Both these shocks follow stationary AR(1) processes. Our construction of the efficient level of output implies that the counterfactual environment in which prices are flexible is a parallel universe, which evolves independently from the outcomes observed in the actual economy (Neiss and Nelson, 2003). In this parallel universe, the intertemporal Euler equation defines the efficient real interest rate as

$$r_t^e = E_t\gamma_{t+1} + E_t\delta_{t+1} - \omega\left(E_t y_{t+1}^e - y_t^e\right).$$

The efficient real interest rate depends positively on the forecastable component of productivity growth and preference shocks and negatively on the unexpected innovations in efficient output.

Turning now to the supply side of the model, the optimal pricing decisions of firms produce a Phillips curve of the form

$$\tilde{\pi}_t = \xi\left(\omega x_t^e + \varphi\tilde{x}_t\right) + \beta E_t\tilde{\pi}_{t+1} + u_t, \quad (3)$$

where

$$\tilde{\pi}_t \equiv \pi_t - \zeta\pi_{t-1}$$

depends on the degree of indexation to past inflation, parametrized by ζ, and u_t is an AR(1) cost-push shock, generated by exogenous fluctuations in desired markups. These fluctuations are the only source of tradeoff between inflation and real activity in this model.

Without markup shocks, the efficient level of aggregate production can be achieved together with price stability (i.e. $\pi_t = 0$), as we can see by substituting $u_t = 0$ and $y_t = y_t^e$, or $x_t^e = 0$, $\forall t$ in equation (3). This is the first best outcome in this economy, since no price needs to change when aggregate inflation is zero, thus eliminating price dispersion across monopolistic producers and the associated distortions in the allocation of resources (Woodford, 2003). When markup shocks are present, on the contrary, the efficient allocation is no longer feasible because the efficient level of *aggregate* output could only be achieved by allowing cost-push shocks to pass-through to inflation entirely, as we can see by solving equation (3) forward with $y_t = y_t^e$ $\forall t$

$$\pi_t = \zeta \pi_{t-1} + \sum_{s=0}^{\infty} \beta^s E_t u_{t+s}.$$

The resulting fluctuations in inflation would then produce an inefficient dispersion of prices and production levels across varieties. At the other extreme of the policy spectrum, perfect inflation stabilization would require cost-push shocks to show-through entirely in deviations of output from its efficient level. Optimal policy, therefore, will distribute the impact of these shocks between output and inflation, as to balance the objectives of price stability and efficient aggregate production.

One implication of this trade-off is that an ex-ante real interest rate, $i_t - E_t \pi_{t+1}$, set to perfectly shadow the efficient rate of return r_t^e, would not be optimal, although the Euler equation (1) implies that such a policy would close the output gap every period and thus achieve the efficient level of aggregate production. This is the main reason for including some feedback from inflation and the output gap even in the interest rate rules that include r_t^e in their intercept, as we do below.[4]

[4]Another reason is that a policy rule of the form $i_t = r_t^e + E_t \pi_{t+1}$ would not deliver the efficient output uniquely, since the nominal interest rate does not respond more than one-to-one to expected inflation (e.g. Clarida et al., 1999).

2.2. Monetary Policy: Baseline Specification

In the baseline policy specification, the central bank sets the nominal interest rate in response to the current inflation rate and the efficient output gap, with a certain degree of inertia

$$i_t = \rho i_{t-1} + (1-\rho)\left(\phi_\pi \pi_t + \phi_x x_t^e\right) + \varepsilon_t^i. \tag{4}$$

Expression (4) represents a natural starting point for our comparative analysis, since it brings the basic ingredients of the empirical literature on interest rate rules into the context of our DSGE framework. Inflation and real activity are the standard arguments of monetary policy rules at least since Taylor (1993), while interest rate inertia typically improves their ability to fit the data, as shown for example by Clarida et al. (2000). We choose the efficient output gap as the baseline policy measure of real economic developments for internal consistency with the rest of our theoretical apparatus. In fact, in our model this gap is both the fundamental driver of inflation, as shown in equation (3), as well as the measure of slack that is relevant for welfare analysis (e.g. Woodford, 2003).

Yet, equation (4) is the "baseline" policy rule only in an expositional sense: it is the first one we consider, for the reasons just described. In economic terms, however, this rule is a priori on par with all the others we evaluate, since the entire model is re-estimated every time a new rule is introduced, as described in the next section. A posteriori, the baseline rule turns out to be a particularly poor choice to close the model, which is one of our results.[5]

3. Inference

We estimate the model laid out in the previous section—and the variants discussed below—with Bayesian methods, as surveyed for example by An and Schorfheide (2007). Bayesian estimation combines prior information on the parameters with the likelihood function of the model to form a posterior density function. We construct the likelihood using the Kalman filter based on the state space representation of the rational expectations solution of

[5]This result is even stronger in the working paper version where the number of estimated rules is much larger than those reported here.

each model under consideration, assuming a prior of zero for the configuration of parameters that imply indeterminacy.

The observation equations are

$$\begin{aligned} \Delta \log GDP_t &= \gamma + y_t - y_{t-1} + \gamma_t \\ \Delta \log PCE_t &= \pi^* + \pi_t \\ FFR_t &= r + \pi^* + i_t, \end{aligned}$$

where Δ is the first difference operator, GDP_t is real GDP, PCE_t is the core PCE deflator (headline ex-food and energy), and FFR_t is the average effective Federal Funds Rate (henceforth FFR), all sampled at a quarterly frequency. The constants in these equations represent the average growth rate of productivity (γ), the long run inflation target (π^*), and the average real interest rate (r). The sample period runs from 1987:Q3 to 2009:Q3, although the main results are not affected by truncating the sample either at 2008:Q4, when the FFR first hit the zero bound, or at 2006Q4, before the eruption of the recent financial crisis. We start the sample with Alan Greenspan's tenure as Fed chairman because, starting with Taylor (1993), there seems to be general agreement that interest rate setting appropriately characterizes U.S. monetary policy during this period.

Columns 2 through 5 of Table 1 report our choice of the priors, which are maintained across all the model specifications we consider. On the demand side, we calibrate the discount factor as $\beta = 0.99$, and impose a loose prior between zero and one on the habit coefficient η, only slightly favoring higher values. These two parameters, together with the average balanced growth rate γ , determine the slope of the Euler equation (1), $\varphi_\gamma^{-1} \equiv \left(1-\eta_\gamma\right)\left(1-\beta\eta_\gamma\right)$, where $\eta_\gamma \equiv \eta e^{-\gamma}$.

On the supply side, the prior on the indexation parameter ζ is centered around 0.6, but is quite dispersed over the unit interval. The slope of the Phillips curve is a convolution of deep parameters, $\xi = (1-\alpha)(1-\alpha\beta)/[\alpha(1+\omega\theta)]$ where α is the fraction of firms that do not change their price in any given period, θ is the elasticity of demand faced by each monopolistic producer and ω is the inverse Frisch elasticity of labor supply. Only the slope ξ can be identified from our observables. We formulate our prior on this parameter as a Gamma

distribution with mean 0.1. This value is somewhat higher than the partial information estimates of the New Keynesian Phillips curve (e.g. Galí and Gertler, 2007; Sbordone, 2002), but is consistent with the low degree of price stickiness found in microeconomic studies such as Bils and Klenow (2004), given reasonable values for ω and θ.[6]

Turning now to the interest rate rule, the prior on the smoothing parameter ρ follows a Beta distribution centered at 0.7, with a 90% probability interval wide enough to encompass most existing estimates. The priors for the feedback coefficients on inflation ϕ_π and real activity ϕ_x are normally distributed with means 1.5 and 0.5 respectively, as in the original Taylor (1993) rule.

The autocorrelations of the exogenous shocks, the ρ_i's in the table (for each shock i), have Beta prior distributions with mean 0.5, while the standard deviations, denoted by σ_i, have Inverse Gamma prior distributions centered at 0.5.

We obtain the posterior mode and inverse Hessian by minimizing the negative of the log posterior density function and use Markov Chain Monte Carlo (MCMC) methods, more specifically a Random Walk Metropolis algorithm, to build a representative sample of the parameters' joint posterior distribution. We monitor the convergence of the chains of draws in each step using a variety of tests. Finally, upon convergence, we combine the chains in the last step, after discarding the initial 25% of the draws in each chain, to form a full sample of the posterior distribution, which represents the source of our inference information.[7]

To evaluate the fit of different policy rules, we compare the marginal data densities (or posterior probabilities) of the corresponding DSGE models. All these models share the tastes and technologies that result in the aggregate demand and supply equations described in Section 2.1, but each is closed with a different interest rate rule. [8]

[6]For example, with $\omega = 1$ and $\theta = 8$, which corresponds to a desired markup of 14%, $\xi = 0.1$ implies $\alpha = 0.4$, or an expected duration of prices of about five months.

[7]Detailed convergence and inference analysis for each specification discussed in the paper is available upon request.

[8]We follow the standard practice (e.g. Lubik and Schorfheide, 2007) that we, as econometricians, remain agnostic about the policy rule while agents in the model exactly know the rule, as well as the rest of the structure of the model, in each case.

We estimate each model separately on the same data, and compute its posterior probability using Geweke (1999)'s modified harmonic mean estimator. To compare fit, we calculate the log of the Bayes factor (multiplied by two) of each alternative specification against the baseline (i.e. the model closed with policy rule (4)). Kass and Raftery (1995) recommend this measure of relative fit since its scale is the same as that of a classic Likelihood Ratio statistic.[9] This procedure results in an overall ranking of the interest rate rules under consideration (as well as in a measure of their individual fit against a common benchmark) and thus implicitly against each other.

3.1. Estimation Results in the Baseline Model

Columns 6 through 9 of Table 1 report selected moments of the marginal posterior distributions of the parameters under the baseline interest rate rule. Although the data are quite informative on most parameters and many of the posterior estimates fall within reasonable ranges, close inspection of the results reveals some puzzling features. To better visualize these anomalies, Figure 1 graphs the prior and posterior marginal distributions for a subset of the parameters.

First, the posterior estimate of the slope of the Phillips curve ξ is minuscule, with a mean of 0.002, two orders of magnitude smaller than the prior mean and at the extreme lower end of the available estimates in the DSGE literature (see for example the survey by Schorfheide, 2008). This posterior estimate implies no discernible trade-off between inflation and real activity, so that inflation is close to an exogenous process driven by movements in desired markups. As a consequence, there is little hope of distinguishing between dynamic inflation indexation and persistent markup shocks as drivers of the observed inflation persistence. This lack of identification is reflected in the bimodal marginal posterior distributions of the parameters ζ and ρ_u—which are generated by MCMC draws with high ζ and low ρ_u or

[9] The Bayes factor of model 1 against model 2 is the ratio of their marginal likelihoods. Kass and Raftery (1995) suggest that values of $2 \log BF$ above 10 can be considered very strong evidence in favor of model 1. Values between 6 and 10 represent strong evidence, between 2 and 6 positive evidence, while values below 2 are "not worth more than a bare mention." We refer to this statistic as the KR criterion.

vice versa—that correspond to local peaks of the joint posterior density of similar heights. Finally, the last two panels of Figure 1 show that the estimated parameters of the interest rate rule imply a strong reaction of policy to the output gap and an extremely weak reaction to inflation, with about half of the posterior draws for ϕ_π below one. These values are puzzling, especially in light of the large literature that has argued that a forceful reaction to inflation has been one of the hallmarks of U.S. monetary policy since the mid-eighties.[10]

The anomalies of the posterior distribution highlighted above reduce the baseline model's marginal data density and contribute to its extremely poor overall fit.[11]

4. Evaluating Alternative Interest Rate Rules

Many aspects of our baseline model could be problematic. In the rest of the paper, we focus on one potential source of these problems, which in our judgement has been largely—and surprisingly—overlooked in the DSGE literature: the specification of the interest rate rule. In particular, the inclusion of the efficient real interest rate in the intercept of the policy rule is the one modification of the baseline rule that provides the most significant and reliable improvement in model fit. This relatively minor adjustment contributes to solving some of the anomalous parameter estimates and identification problems illustrated in Figure 1.[12] If we further modify the policy rule to include a time varying inflation target we obtain an even better model fit, but the gain is not as significant as that of introducing the efficient interest rate. Combining both modifications offers the best fit. Furthermore, the resulting improvement in fit is robust to a large set of changes in the other elements of the policy rule, as well as to changes in the model of private sector behavior.

[10] Still, values of ϕ_π lower than one do not necessarily generate indeterminacy, especially for high values of ϕ_x. Following Woodford (2003), in our model the equilibrium is determinate if and only if $\phi_\pi + (1-\beta)\phi_x/\xi > 1$. At the posterior mode, the left-hand side of the previous inequality is slightly smaller than 4.

[11] For a more exhaustive comparison of alternative policy rules, see Cúrdia et al. (2011b). The baseline specification ranks 47th in terms of marginal likelihood, among the 55 evaluated in that version of the paper.

[12] We do not address the extent to which different policy rules aid or hinder the identification of the model's parameters, although this issue would deserve further scrutiny. For a recent study of identification in DSGEs, see Canova and Sala (2009), who find that identification is often problematic in this class of models.

4.1. Tracking the Efficient Real Interest Rate

The idea that an "equilibrium" interest rate (EIR) might represent a useful reference point for monetary policy was familiar to Federal Reserve policymakers well before Woodford (2003) revitalized its Wicksellian roots. For example, in his Humphrey Hawkins testimony to Congress in May 1993, Chairman Alan Greenspan stated that "...In assessing real rates, the central issue is their relationship to an equilibrium interest rate, specifically, the real rate level that, if maintained, would keep the economy at its production potential over time. Rates persisting above that level, history tells us, tend to be associated with slack, disinflation, and economic stagnation—below that level with eventual resource bottlenecks and rising inflation, which ultimately engenders economic contraction. Maintaining the real rate around its equilibrium level should have a stabilizing effect on the economy, directing production toward its long-term potential" (Greenspan, 1993).[13]

Greenspan's quote reflects a shift of emphasis at the Federal Reserve from monetary aggregates to the EIR to gauge the stance of monetary policy that occurred between the end of the 1980s and the beginning of the 1990s (McCallum and Nelson, 2011). Early on, the EIR presented to policymakers was simply the equilibrium real interest rate prevailing in a real business cycle model. As the literature started to incorporate nominal rigidities in the basic neo-classical framework, the EIR became a common feature also of New Keynesian models used to discuss monetary policy issues (see, for example, King and Wolman, 1999, and the survey in Amato, 2005).

In this section, we investigate the extent to which Chairman Greenspan's reasoning had a measurable impact on the evolution of the observed nominal interest rate over our sample. To measure the EIR within our DSGE model, we follow the Chairman's description and compute the counterfactual "real rate level that, if maintained, would keep the economy at its production potential over time." If we define "potential" output as the efficient aggregate

[13]Quantitative measures of the EIR are today a regular input in the monetary policy debate at the Federal Reserve, as demonstrated by the fact that a chart with a range of estimates of the EIR is included in most published Bluebooks at least since May 2001 (see http://www.federalreserve.gov/monetarypolicy/fomc_historical.htm).

level of production y_t^e in the model—as, for instance, in Justiniano et al. (2011)—the EIR corresponds to the efficient rate of return r_t^e. This variable is our preferred measure of the EIR, since its concept is grounded in the microeconomic structure of the DSGE model.[14]

We embed this measure of the EIR within the baseline policy rule, which becomes

$$i_t = \rho i_{t-1} + (1-\rho)\left[r_t^e + \phi_\pi \pi_t + \phi_x x_t^e\right] + \varepsilon_t^i. \tag{5}$$

This specification (which we call Re) improves the model's marginal likelihood by approximately 20 log-points with respect to the baseline, which represents very strong evidence in favor of policy rule (5) according to the aforementioned KR criterion. Part of the overall improvement in fit is due to the increase in explanatory power of the systematic component of the monetary policy rule. The introduction of the EIR in the interest rate rule reduces the standard deviation of the monetary policy shock by more than a third. To the best of our knowledge, this paper is the first to document the good empirical performance of a policy rule that allows for a gradual adjustment of the nominal interest rate to movements in the efficient real rate.

Figure 2 plots the posterior median of r_t^e (blue solid line) and the 90% uncertainty bands (green shaded area) implied by the model under rule (5). The vertical grey areas correspond to the National Bureau of Economic Research recession dates. Overall, the estimated efficient real interest rate is a good business cycle indicator over our sample, rising during booms and dropping sharply in recessions. In the last two recession episodes in our sample, the efficient real interest rate conveys early signals of a future slowdown, decreasing a few quarters before the recession actually starts.[15]

Figure 2 also shows that the inferred movements in r_t^e mirror quite closely those in the FFR (red dashed line), which helps to further explain the empirical success of equation

[14] Cúrdia et al. (2011b) also reports results for alternative definitions of the EIR, corresponding to the potential output implied by the statistical filters described in section 4.2.1. None of these alternatives improves on the specification presented here.

[15] In the 1990 recession, the drop of r_t^e coincides with the beginning of the recession. Interestingly, the EIR behaves very similarly to more standard measures of slack in the real economy, such as output gaps. See Cúrdia et al. (2011b) for more details.

(5) as a description of monetary policy. The close co-movement between the FFR and the estimates of r_t^e may, however, raise the concern that the observations on the nominal interest rate fully "explain" the estimates of r_t^e, and not vice versa. Figure 3 shows that this is clearly not the case. The estimated time path of r_t^e when the interest rate rule includes the EIR (thick green solid line) is almost identical to the baseline case (thin blue solid line). This result is also robust to the introduction of a time-varying inflation target (red dashed line), that we discuss in the next subsection. The main difference between the estimates (not shown in figure 3) is that the posterior distribution is tighter when r_t^e enters the interest rate rule. This enhanced precision of the estimates suggests that the nominal interest rate carries useful information on r_t^e in the specification Re, although this information does not distort the inference on its average time-path.

Some intuition for the robustness of the estimates of r_t^e across models can be gleaned from the expression for the efficient real interest rate derived in section 2, which we report here for convenience

$$r_t^e = E_t\gamma_{t+1} + E_t\delta_{t+1} - \omega\left(E_t y_{t+1}^e - y_t^e\right).$$

If the log-deviations of efficient output from the balanced growth path were a martingale (i.e. $E_t y_{t+1}^e = y_t^e$), the efficient real interest rate would be equal to the sum of the forecastable movements in the growth rate of productivity γ_t and in the intertemporal taste shock δ_t. In our estimated models, the deviations from the condition $E_t y_{t+1}^e = y_t^e$ are "small,"as are the forecastable movements in γ_t. The taste shock δ_t, on the contrary, is persistent, and its innovations are sizable, so that its forecastable movements tend to be the main driving force of movements in r_t^e. Our estimates precisely and robustly pin down the cyclical behavior of these forecastable movements in δ_t, with little variation across specifications. As a result, the inference on the evolution of the efficient real rate over time is remarkably consistent across all the models we consider.

The key role of the intertemporal shock δ_t in reconciling this class of DSGE models with the data is a manifestation of the well-known deficiencies of standard Euler equations in pricing returns, as first documented by Hansen and Singleton (1982) and more recently re-emphasized in a DSGE context by Primiceri et al. (2006). According to this interpretation,

intertemporal shocks are reduced-form measures of exogenous changes in financial factors. The tight connection between r_t^e and intertemporal shocks and the fact that its inclusion in the interest rate rule significantly improves the fit of our DSGE model represent suggestive evidence of the importance of financial conditions for the conduct of monetary policy in the last two decades.[16]

4.2. Robustness to Changes in the Policy Rule

The evidence presented in the previous section suggests that the efficient real interest rate r_t^e is a potentially useful indicator of the stance of monetary policy, since its inclusion among the arguments of the baseline policy rule significantly improves the ability of our model to fit the data. This section offers some evidence that this result does not depend on the arbitrary choice of the baseline specification. Regardless of how we specify the other arguments of the policy rule, the inclusion of r_t^e uniformly improves the model's empirical performance. Aside from illustrating the robustness of the paper's main result, this exercise allows us to explore the empirical fit of other interest rate rules. We consider two broad classes of alternative rules. First, we experiment with a different measure of the output gap. Second, we allow for time-variation in the inflation target π^*. Some of these rules further improve on the empirical performance of (5).

Table 2 reports the overall fit of several alternative rules. The column labeled "logML" shows the log-marginal likelihood for each specification, while the column labeled "KR" shows the Kass and Raftery (KR) measure described in Section 3 (where positive values imply an improvement in fit relative to the "Baseline" specification). The rules "Baseline" and "Re" correspond to the specifications (4) and (5), respectively. Both of these specifications have been discussed above. The next subsections discuss the other formulations.

[16] Shocks to r_t^e are sometimes used as a shortcut for more fundamental financial disturbances that generate large crises (e.g. Eggertsson, 2008). In the early stages of the recent financial crisis, McCulley and Toloui (2008), Meyer and Sack (2008), and Taylor (2008) proposed spread-adjusted Taylor rules to appropriately capture the effects of the turmoil in credit markets on the level of the FFR. In a model in which credit spreads induce fluctuations in the efficient real interest rate, Cúrdia and Woodford (2010) show that such an adjustment is indeed desirable, although less than to a one-to-one extent.

4.2.1. Output Gap

The measure of economic slack that we chose to include in the baseline interest rate rule (4) is the deviation of real GDP from its efficient level. This choice is fairly common in DSGE work (e.g. Smets and Wouters, 2007), although far from universal. One drawback of this approach is that the resulting policy rule is impossible to compare with those estimated in the vast literature that employs partial information econometric techniques (see Coibion and Gorodnichenko (2011) for a recent contribution and survey of this literature), since the construction of the counterfactual efficient level of output requires a general equilibrium model. Moreover, the efficient output gap might be considered an implausible choice as a summary statistic for policymakers' views on the level of resource utilization, precisely because of its model dependency.

To bridge the gap between our general equilibrium framework and the work based on single equation methods, we consider the Hodrick and Prescott (HP) filter as a tool to construct smooth versions of potential output, given its popularity in applied macroeconomics.[17]

One difficulty in making the HP filter operational within a DSGE model is that its ideal representation is a two-sided, infinite moving average, whose standard approximation to finite samples requires different coefficients on the observations at the beginning, in the middle, and at the end of the sample. Such a pattern of coefficients is difficult to replicate within a dynamic system of rational expectation equations with a parsimonious state space. To circumvent this problem, we adapt the methodology proposed by Christiano and Fitzgerald (2003) for the approximation of ideal band pass filters. Christiano and Fitzgerald (2003) suggest to use forecasts (and backcasts) from an auxiliary time-series model—in their case a simple unit root process—to extend the sample in the past and in the future. In our implementation of their idea, the auxiliary model that generates the dummy observations is the linearized DSGE itself.

This approach is particularly convenient for our purposes because it produces a very

[17]See Orphanides and Van Norden (2002) for a comprehensive survey of the use of statistical filters as measures of the output gap and their pitfalls.

parsimonious recursive expression for the DSGE-HP gap

$$\left[1+\lambda(1-L)^2\,(1-F)^2\right]x_t^{HP}=\lambda(1-L)^2\,(1-F)^2\,y_t, \tag{6}$$

where the operators L and F are defined by $Ly_t = y_{t-1}$ and $Fy_t = E_t y_{t+1}$, and the smoothing parameter λ is set at the typical quarterly value of 1600. This expression can thus be added to the system of rational expectations equations that defines the equilibrium of the model. More details on the derivation of equation 6 and on its interpretation, together with some background on linear filtering, can be found in Appendix B.

The time series for the output gap obtained through this procedure (DSGE-HP) is very similar to the standard finite sample approximation of the HP filter. This result is comforting and supports our use of the DSGE-HP filter as an effective de-trending tool, which produces a measure of capacity utilization similar to those often used in single-equation estimates of the Taylor rule.

When we estimate the model replacing the efficient output gap with x_t^{HP} in the interest rate rule, the presence of r_t^e continues to make a difference, although the improvement becomes less significant. Without the efficient real interest rate in the policy rule, the KR criterion deteriorates by four points, as shown in Table 2. To be sure, both specifications perform substantially better than the baseline but the presence of the EIR reduces the need for a measure of slack to be included in the policy rule. In fact, a rule with the efficient real interest rate but no output gap (i.e. with $\phi_x = 0$) ranks slightly better than a rule with the efficient real interest rate and the statistical output gap x_t^{HP}.

Overall, conditional on the presence of the efficient real interest rate, the rule with the model-based efficient output gap ranks higher than any rule with a statistical filter of output.[18]

4.2.2. Time-Varying Inflation Target

In this section, we further enlarge the set of policy rules subject to our evaluation, by introducing a time-varying inflation target (TVIT), a fairly common feature in the recent

[18] Cúrdia et al. (2011b) consider additional statistical filters, including smoother trends and exponential filters. Also for these rules, the presence of r_t^e is crucial.

empirical DSGE literature (Ireland, 2007; Cogley et al., 2010). This addition creates a new class of feedback rules, of the form

$$i_t = \rho i_{t-1} + (1-\rho)\left[r_t^e + \pi_t^* + \phi_\pi\left(\pi_t - \pi_t^*\right) + \phi_x x_t\right] + \varepsilon_t^i, \tag{7}$$

where π_t^* is an exogenous AR(1) process that represents persistent deviations of the inflation target from its long-run value π^*.[19]

Introducing a TVIT in the baseline policy rule helps capture the low-frequency movements in inflation and the nominal interest rate that are evident even in our relatively short sample. Inflation hovered around 4% in the late 1980s and started to converge toward its more recent range around 2% only after the recession of the early 1990s. This process of so-called "opportunistic disinflation" took until the middle of the decade to complete. One simple way of capturing the central bank's willingness to delay the achievement of its ultimate inflation objective until the "next" recession, which is at the heart of the opportunistic approach to disinflation, is to allow smooth time-variation in its short-run inflation target, as in specification (7).

The fit of the model improves by 24 points in the KR scale (rule "Pistar" in Table 2), very strong evidence in favor of the inclusion of a TVIT in the policy rule. However, this improvement is still not nearly as good as the introduction of the EIR by itself in the rule. The best rule is the one that combines these two components, denoted by "RePistar," which achieves an improvement over the baseline case of 54 points in the KR scale. Adding the EIR to the Pistar specification improves the KR criterion by 30 points, which constitutes a significant increase in fit. Interestingly, this gain is comparable to the one obtained when the EIR is included in the equivalent specifications with a constant inflation target, i.e. when comparing the baseline rule to rule Re.

Table 3 suggests a similar conclusion by looking at the fit of the interest rate rule only. The first column ("In-sample") reports the posterior median standard error of the monetary policy shock ε_t^i over the sample.[20] The second column of the table ("Population") is the

[19] The autocorrelation coefficient of π_t^* has a Beta prior tightly distributed around a mean of 0.95.

[20] We draw from the posterior distributions of the parameters 1000 times. For each draw, we compute the

posterior median of the estimated standard deviation of the monetary policy shock σ_i. As mentioned for the EIR, the introduction of a TVIT reduces the standard error of the monetary policy rule by roughly one third. The combination of both elements lowers this measure even further, although the improvement is not as big as for the KR criterion.

These results suggest that the efficient real rate and a smoothly evolving inflation target enhance the empirical performance of the model through fairly independent channels and should thus be complementary features in policy specifications with good empirical properties.

As previously anticipated, figure 3 shows that the introduction of the TVIT does not substantially change the behavior of the posterior estimates of the EIR (red dashed line). This measure remains fairly resilient to changes in the policy rule. The reason why the EIR and the TVIT are complementary features of the best specification is that the former helps improve the business cycle properties of the model, while the latter helps capture the low frequency component of inflation.

The overall improvement in fit also reflects in better parameter estimates. Figure 4 compares the marginal prior and posterior distributions for selected model parameters. The estimated slope of the Phillips curve is now 0.018 at the posterior median, compared to less than 0.002 under the baseline rule. Given the calibrated parameters, the implied average frequency of price adjustment when evaluating ξ at the posterior median is a much more reasonable 4 quarters, compared to more than 8 in the baseline case.

Furthermore, the posterior distributions for the indexation parameter ζ and the persistence of the markup shock ρ_u are no longer bimodal. The disappearance of the multiple modes depends on the introduction of the EIR in the policy rule. The posterior distributions for these two parameters in the Re specification also display single modes. The presence of the TVIT, however, plays a key role for the persistence of the markup shock. In its absence, the posterior distribution of ρ_u concentrates on high values. Figure 4 shows that the opposite

standard deviation of the difference between the observed FFR and the FFR predicted by the systematic component of each monetary policy rule. We report the median of the distribution of standard errors obtained.

result holds in the RePistar specification. In the absence of a TVIT, the model attributes the low frequency movements in inflation to the only source of persistence available, which is the markup shock process. When we introduce the TVIT, the data favor this channel to explain the low frequency movements in inflation.

Finally, the posterior distribution for the feedback coefficients on inflation ϕ_π and the output gap ϕ_x are now more in line with the conventional values discussed in the literature. If anything, the posterior distribution ϕ_x is not very different from the prior, suggesting possible lack of identification. Figure 5 hints to the potential reason. This picture shows the posterior median estimate of the output gap measure under three different specifications for the interest rate rule (Baseline, Re, RePistar). Unlike the EIR, the posterior estimates of the output gap significantly depend on the policy rule specification. In the baseline specification (thin blue solid line), the output gap captures the business cycle reasonably well. However, when the policy rule includes the EIR only (thick green solid line), the output gap partly looses its cyclical properties, especially early in the sample. The output gap in the RePistar specification (red dashed line), albeit much more muted than in the baseline case, retains some features of a business cycle indicator but the uncertainty surrounding the posterior estimates (not shown) casts doubts on the usefulness of this measure even in this case.

Taken together, these results suggest that, to appropriately account for the movements in the FFR since the Greenspan's tenure, the interest rate rule that closes a monetary DSGE model should feature the efficient interest rate and a time-varying inflation target. Other measures of slack in real activity, such as the output gap, become almost completely redundant. Indeed, Table 2 shows that the fit of the model does not dramatically worsen if we consider an interest rate rule with EIR and TVIT but no output gap.

5. A Medium-Scale DSGE Model

We conclude our investigation of interest rate rules by evaluating the robustness of the results obtained so far within a medium-scale DSGE model, along the lines of Christiano et al. (2005) and Smets and Wouters (2007). The exact specification we adopt for the private sector behavior follows the work of Justiniano et al. (2010, henceforth, JPT), to

which we refer the reader for further details. Within this framework, we revisit the role of the equilibrium interest rate and time-varying inflation target in the policy rule. In order to properly compare the results to those of the small model, we estimate the JPT model on the same same set of observables—GDP growth, inflation and the FFR—and on the same sample. This way any changes in relative ranking are solely due to the increased complexity of the model.

Table 4 ranks the improvement in fit of changing the policy rule from the baseline to one that includes the EIR (Re), the TVIT (Pistar), or both (RePistar). For ease of comparison, we report the KR measure for both the small model previously considered and for the JPT model.[21]

The results strongly corroborate the findings from the small model. First, adding a time-varying inflation target to the rule increases the fit by similar magnitudes. Second, the presence of the equilibrium interest rate in the interest rate rule increases the fit even more. Interestingly, the increase in fit in the JPT model is even larger than in the small model. Third, adding both the EIT and TVIT to the policy rule achieves the best fit among these alternatives—and the increase in fit is again higher than in the small model.

Beyond fit, one important question is whether a more complex model structure affects the estimated path for the EIR. Figure 6 shows the posterior median for this variable in the small model (dark/blue) and in JPT (light/red) for the specifications with (thick dashed line) and without the TVIT (thin solid line). The first observation is that the presence of the TVIT does not make a substantial difference for the estimated path of the EIR. Qualitatively, the different models produce remarkably similar measures of the EIR. In particular, r_t^e always drops at beginning of each recessions in our sample, or just before. The main difference between small model and JPT is quantitative. The estimated path of the EIR is less volatile in JPT. Consequently, the estimated depth of contractions in the JPT model is milder, especially for the last two recession episodes, and booms are more moderate. Yet the main

[21] The interested reader may want to refer to Cúrdia et al. (2011b) for additional rules estimated in the JPT model.

result stands: the EIR is a good business cycle indicator in both the small model and JPT, independently of the presence of a TVIT.

In sum, this section confirms that adding the EIR to the policy rule increases the fit of the model with respect to inflation, output and the interest rate. Furthermore, the posterior estimates of the EIR do not significantly depend on the presence of a TVIT. Finally, the EIR remains quite comparable across models of different degree of complexity.

6. Conclusions

The positive DSGE literature focuses an overwhelming share of its modeling attention on the behavior of the private sector, while treating that of the central bank as an afterthought. This state of affairs is not too surprising, since reducing the complexity of the private sector to fit into a macroeconomic model offers a vast menu of modeling choices. In comparison, capturing the behavior of one monetary authority is certainly easier and perhaps less controversial. Yet, paying virtually no attention to this step in the specification of a general equilibrium model seems suboptimal, for at least two reasons. First, in the current vintage of monetary DSGE models, the systematic response of the central bank to economic developments can have significant effects on the equilibrium, as demonstrated by the vast body of normative work in the field (see Woodford (2011) for a survey). Second, one of the main objectives of DSGE models is to offer a quantitative tool to study the consequences of different approaches to the conduct of monetary policy. This study is complicated by the lack of systematic guidance on the extent to which different plausible policy rules, once embedded into a general equilibrium apparatus, contribute to its ability to account for the historical relations between the macroeconomic variables of interest.

This paper—and the research that it summarizes—attempted to provide some of that guidance by estimating a large set of interest rate rules in the context of two standard DSGE models and comparing their empirical fit. The most notable (and robust) result of this extensive investigation is that allowing the interest rate to track the evolution of the efficient real interest rate—the real interest rate that would prevail in equilibrium if the economy were perfectly competitive—significantly improves the fit of the DSGE models we

estimated. In fact, this measure of the equilibrium interest rate is a better proxy for the real economic developments to which monetary policy seems to respond than any of the several measures of the output gap we experimented with.

Our results are subject to two important caveats. First, model specification matters, since our criterion of fit depends on the way the policy rule interacts with the rest of the model. This is a feature of the general equilibrium nature of the models we work with, and of the full information empirical methods we adopted for their estimation. These methods are necessary for our purposes, however, since the efficient real interest rate is a counterfactual equilibrium object built within the model. More work across different models would therefore be desirable, although we attempted to address this issue by illustrating the robustness of our results across two fairly popular DSGE specifications. Second, model comparison through marginal data densities and Bayes factors applied to DSGE models is subject to some pitfalls, highlighted for example by Del Negro and Schorfheide (2011). However, the large differences in fit we uncovered suggest that the specification of the policy rule does make a difference.

Going forward, we expect to devote some of our research to further scrutinize the role of the efficient real interest rate r_t^e as a useful explanatory factor for the movements in observed nominal interest rates. In particular, we would like to better understand the origins of this combination of shocks, which especially in our baseline model is partly a reflection of the empirical shortcomings of the intertemporal Euler equation, as captured by the presence of the shock δ_t. Moreover, it would be interesting to explore more realistic assumptions on the information available to policy makers when making their decisions, focusing in particular on the fact that the efficient real interest rate is not observable in practice, unlike in our model. In a similar vein, the use of real-time data would provide an interesting further perspective on our results.

References

Adolfson, M., Laseen, S., Linde, J., Svensson, L. E., 2008. Optimal monetary policy in an operational medium-sized dsge model. NBER Working Paper No. 14092.

Amato, J., 2005. The role of the natural rate of interest in monetary policy. BIS Working Paper No. 171.

An, S., Schorfheide, F., 2007. Bayesian analysis of dsge models. Econometric Reviews 26, 113–172.

Baxter, M., King, R., 1999. Measuring business cycles: Approximate band-pass filters for economic time series. Review of Economics and Statistics 81, 575–593.

Bils, M., Klenow, P., 2004. Some evidence on the importance of sticky prices. Journal of Political Economy 112, 947–985.

Blanchard, O., Fischer, S., 1989. Lectures on Macroeconomics. MIT Press.

Calvo, G., 1983. Staggered prices in a utility-maximizing framework. Journal of Monetary Economics 12 (3), 383–398.

Canova, F., Sala, L., 2009. Back to square one: Identification issues in dsge models. Jounral of Monetary Economics 56, 431–449.

Christiano, L., Eichenbaum, M., Evans, C., February 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113 (1), 1–45.

Christiano, L., Fitzgerald, T., 2003. The band pass filter. International Economic Review 44, 435–465.

Clarida, R., Galí, J., Gertler, M., 1999. The science of monetary policy: A new keynesian perspective. Journal of Economic Literature 37, 1661–1707.

Clarida, R., Galí, J., Gertler, M., February 2000. Monetary policy rules and macroeconomic stability: Evidence and some theory. The Quarterly Journal of Economics 115 (1), 147–180.

Cogley, T., Primiceri, G. E., Sargent, T., 2010. Inflation-gap persistence in the u.s. American Economic Journal: Macroeconomics 2, 43–69.

Coibion, O., Gorodnichenko, Y., 2011. Monetary policy, trend inflation, and the great moderation: An alternative interpretation. American Economic Review 101, 341–370.

Cúrdia, V., Del Negro, M., Greenwald, D., 2011a. Rare large shocks in the u.s. business cycle.

Cúrdia, V., Ferrero, A., Ng, G., Tambalotti, A., 2011b. Evaluating interest rate rules in an estimated dsge model. Federal Reserve Bank of New York Staff Report 510.

Cúrdia, V., Woodford, M., 2010. Credit spreads and monetary policy. Journal of Money Credit and Banking 42, 3–35.

Del Negro, M., Schorfheide, F., 2011. Bayesian macroeconometrics. In: Geweke, J., Koop, G., Van Dijk, H. (Eds.), Handbook of Bayesian Econometrics. Oxford University Press, pp. 293–289.

Eggertsson, G., 2008. Great expectations and the end of the depression. American Economic Review 98, 1476–1516.

English, W., Nelson, W., Sack, B., 2003. Interpreting the significance of the lagged interest rate in estimated monetary policy rules. Contributions to Macroeconomics 3, Article 5.

Galí, J., Gertler, M., Fall 2007. Macroeconomic modeling for monetary policy evaluation. Journal of Economic Perspectives 21 (4), 25–46.

Geweke, J., 1999. Using simulation methods for bayesian econometric models: Inference, development,and communication. Econometric Reviews 18, 1–126.

Greenspan, A., 1993. Statement on the conduct of monetary policy before the subcommitte on economic growth and credit formation of the committee on banking, finance and urban affairs of the house of representatives.

Hansen, L. P., Singleton, K., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. Econometrica 50, 1269–1286.

Hodrick, R., Prescott, E., 1997. Post-war u.s. business cycles: An empirical investigation. Journal of Money, Credit, and Banking 29, 1–16.

Ireland, P., 2007. Changes in the federal reserve?s in?ation target: Causes and consequences. Journal of Money, Credit, and Banking 39, 1851–1882.

Judd, J. P., Rudebusch, G. D., 1998. Taylor's rule and the fed: 1970-1997. Federal Reserve Bank of San Francisco Economic Review 3, 3–16.

Juillard, M., Kamenik, O., Kumhof, M., Laxton, D., 2006. Measures of potential output from an estimated dsge model of the united states. Czech National Bank Working Paper 2006/11.

Justiniano, A., Primiceri, G., Tambalotti, A., 2010. Investment shocks and business cycles. Jounral of Monetary Economics 57, 132–145.

Justiniano, A., Primiceri, G., Tambalotti, A., 2011. Measures of potential output from an estimated dsge model of the united states.

Kass, R., Raftery, A., 1995. Bayes factors. Journal of the American Statistical Association 90, 773–795.

King, R., Rebelo, S., 1993. Low frequency filtering and real business cycles. Journal of Economic Dynamics and Control 17, 201–231.

King, R., Wolman, A., 1999. What should the monetary authority do when prices are sticky? In: Taylor, J. (Ed.), Monetary Policy Rules. University of Chicago Press, pp. 349–404.

Lubik, T. A., Schorfheide, F., 2007. Do central banks respond to exchange rate movements? a structural investigation. Journal of Monetary Economics 54, 1069–1087.

McCallum, B., Nelson, E., 2011. Money and inflation: Some critical issues. In: Friedman, B. M., Mic (Eds.), Handbook of Monetary Economics. Vol. 3A. Elsevier/North-Holland, pp. 97–153.

McCulley, P., Toloui, R., February 2008. Chasing the neutral rate down: Financial conditions, monetary policy, and the taylor rule. Global central bank focus, PIMCO.

Meyer, L. H., Sack, B. P., March 2008. Updated monetary policy rules: Why don't they explain recent monetary policy? Monetary policy insights, Macroeconomic Advisors.

Neiss, K., Nelson, E., 2003. The real-interest-rate gap as an inflation indicator. Macroeconomic Dynamics 7, 239–262.

Orphanides, A., 2003. Historical monetary policy analysis and the taylor rule. Journal of Monetary Economics 50, 983–1022.

Orphanides, A., Van Norden, S., 2002. The unreliability of output-gap estimates in real time. Review of Economics and Statistics 84, 569–583.

Primiceri, G., Schaumburg, E., Tambalotti, A., 2006. Intertemporal disturbances. NBER Working Paper No. 12243.

Sbordone, A., 2002. Prices and unit labor costs: A new test of price stickiness. Journal of Monetary Economics 49, 265–292.

Schmitt-Grohé, S., Uribe, M., 2007. Optimal simple and implementable monetary and fiscal rules. Journal of Monetary Economics 54, 1702–1725.

Schorfheide, F., 2008. Dsge model-based estimation of the new keynesian phillips curve. Federal Reserve Bank of Richmond Economic Quarterly 94, 397–433.

Smets, F., Wouters, R., 2007. Shocks and frictions in us business cycles: A bayesian dsge approach. American Economic Review 97 (3), 586–606.

Svensson, L. E., 2003. What is wrong with taylor rules? using judgment in monetary policy through targeting rules. Journal of Economic Literature 41, 426–477.

Taylor, J. B., 1993. Discretion versus policy rules in practice. Carnegie-Rochester Conference Series on Public Policy 39, 195–214.

Taylor, J. B., 2008. Monetary policy and the state of the economy. Testimony before the Committee on Financial Services, U.S. House of Representatives, February 26, 2008.

Trehan, B., Wu, T., 2007. Time-varying equilibrium real rates and monetary policy analysis. Journal of Economic Dynamics and Control 31, 1584–1609.

Watson, M., 2007. How accurate are real-time estimates of output trends and gaps? Federal Reserve Bank of Richmond Economic Quarterly 93, 143–161.

Woodford, M., 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press.

Woodford, M., 2011. Optimal monetary stabilization policy. In: Friedman, B. M., Woodford, M. (Eds.), Handbook of Monetary Economics. Vol. 3B. Elsevier/North-Holland, Amsterdam, pp. 723–828.

Parameter	Prior				Posterior			
	Distribution	5%	Median	95%	Mode	5%	Median	95%
ω	$G(1, 0.2)$	0.70	0.99	1.35	0.96	0.67	0.96	1.32
ξ	$G(0.1, 0.05)$	0.03	0.09	0.19	0.001	0.000	0.002	0.004
η	$B(0.6, 0.2)$	0.25	0.61	0.90	0.56	0.46	0.59	0.70
ζ	$B(0.6, 0.2)$	0.25	0.61	0.90	0.69	0.10	0.53	0.80
ρ	$B(0.7, 0.15)$	0.43	0.72	0.92	0.70	0.62	0.72	0.80
ϕ_π	$N(1.5, 0.25)$	1.09	1.50	1.91	0.89	0.66	1.03	1.49
$4\phi_x$	$N(0.5, 0.2)$	0.17	0.50	0.83	1.19	0.97	1.21	1.45
$400\pi^*$	$N(2, 1)$	0.36	2.00	3.64	2.36	1.89	2.38	2.85
$400r^a$	$N(2, 1)$	0.36	2.00	3.64	1.90	0.83	1.90	2.95
$400\gamma^a$	$N(3, 0.35)$	2.42	3.00	3.58	2.94	2.48	2.94	3.40
ρ_δ	$B(0.5, 0.2)$	0.17	0.50	0.83	0.92	0.87	0.92	0.95
ρ_γ	$B(0.5, 0.2)$	0.17	0.50	0.83	0.56	0.29	0.53	0.72
ρ_u	$B(0.5, 0.2)$	0.17	0.50	0.83	0.11	0.06	0.32	0.71
σ_δ	$IG1(0.5, 2)$	0.17	0.34	1.24	1.23	0.95	1.37	2.04
σ_γ	$IG1(0.5, 2)$	0.17	0.34	1.24	2.05	1.34	2.13	3.02
σ_u	$IG1(0.5, 2)$	0.17	0.34	1.24	0.53	0.19	0.43	0.60
σ_i	$IG1(0.5, 2)$	0.17	0.34	1.24	0.27	0.23	0.30	0.39

Table 1: Prior and posterior marginal distributions for the paramaters in the baseline model. G stands for Gamma, B stands for Beta, N stands for Normal and IG1 stands for Inverse Gamma 1, with mean and standard deviation in parenthesis.

	without r^e			with r^e	
RULE	logML	KR	RULE	logML	KR
Baseline ($\phi_\pi \pi_t + \phi_x x_t^e$)	-379.0	0.0	Re	-359.3	39.4
Pistar	-366.8	24.4	RePistar	-351.8	54.4
HP ($\phi_\pi \pi_t + \phi_x x_t^{HP}$)	-368.2	21.6	ReHP	-366.1	25.7
PistarHP	-361.6	34.9	RePistarHP	-356.2	45.6
NoGap ($\phi_\pi \pi_t$)	-393.2	-28.3	ReNoGap	-364.6	28.9
PistarNoGap	-371.4	15.1	RePistarNoGap	-353.6	50.7

Table 2: Small model rules. The first column presents the specification with the base rule, log marginal likelihood and KR criterion. The second column presents the corollary rule (*Re*) that includes a response to the efficient real interest rate r^e. Rules with *Pistar* in the name include a time-varying inflation target π^* and deviations from the target in the rule as discussed in Section 4.5, e.g. in the form of $i_t = \rho i_{t-1} + (1 - \rho)[r_t^* + \pi_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_x x_t] + \epsilon_t^i$. Rules with *HP* in the name use a measure of the output gap discussed in (6). Rules with *NoGap* in the name mean that the interest rate does not respond directly to the output gap ($\phi_x = 0$).

Rule	In-Sample	Population
Baseline: $\phi_\pi \pi_t + \phi_x x_t^e$	0.2982	0.3001
Pistar: $\pi_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_x x_t^e$	0.2048	0.2088
Re: $r_t^e + \phi_\pi \pi_t + \phi_x x_t^e$	0.1907	0.1940
RePistar: $r_t^e + \pi_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_x x_t^e$	0.1793	0.1827

Table 3: In-Sample and Population estimates of the standard deviation of the monetary policy shock for the four major rules in the Small model. The In-Sample estimate is the posterior median of the standard error of the monetary policy rule equation residuals from the estimation period. The Population estimate is the posterior median of the estimated standard deviation on the monetary policy shock σ_i.

Rule	Small Model	JPT
Baseline: $\phi_\pi \pi_t + \phi_x x_t^e$	0.0	0.0
Pistar: $\pi_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_x x_t^e$	24.4	25.7
Re: $r_t^e + \phi_\pi \pi_t + \phi_x x_t^e$	39.4	54.6
RePistar: $r_t^e + \pi_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_x x_t^e$	54.4	65.9

Table 4: Kass and Raftery (KR) criterion for the four major rules against their respective Baseline specifications for the small and JPT models.

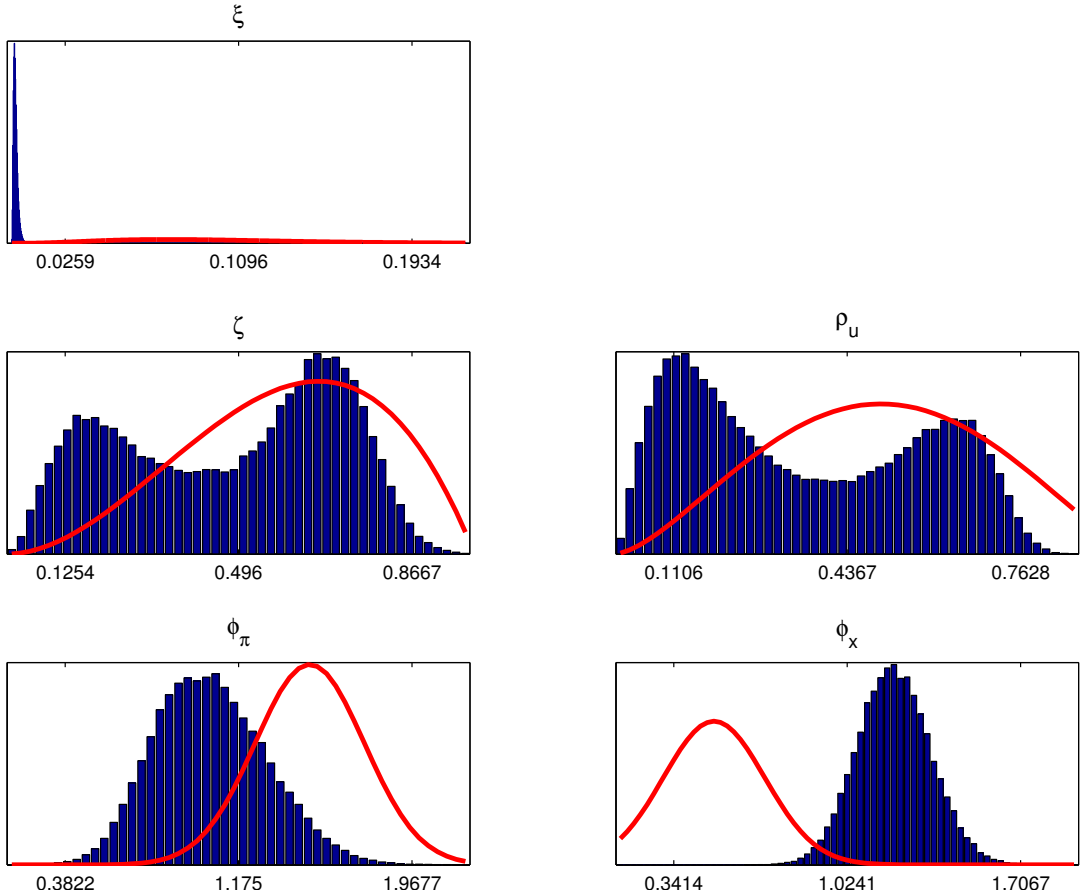

Figure 1: Prior and posterior distributions for ξ, ζ, ρ_u, ϕ_π, and ϕ_x under the baseline specification of interest rate rule: $i_t = \rho i_{t-1} + (1-\rho)(\phi_\pi \pi_t + \phi_x x_t^e) + \epsilon_t^i$. For each parameter, the solid red line represents the prior while the blue histogram is the posterior.

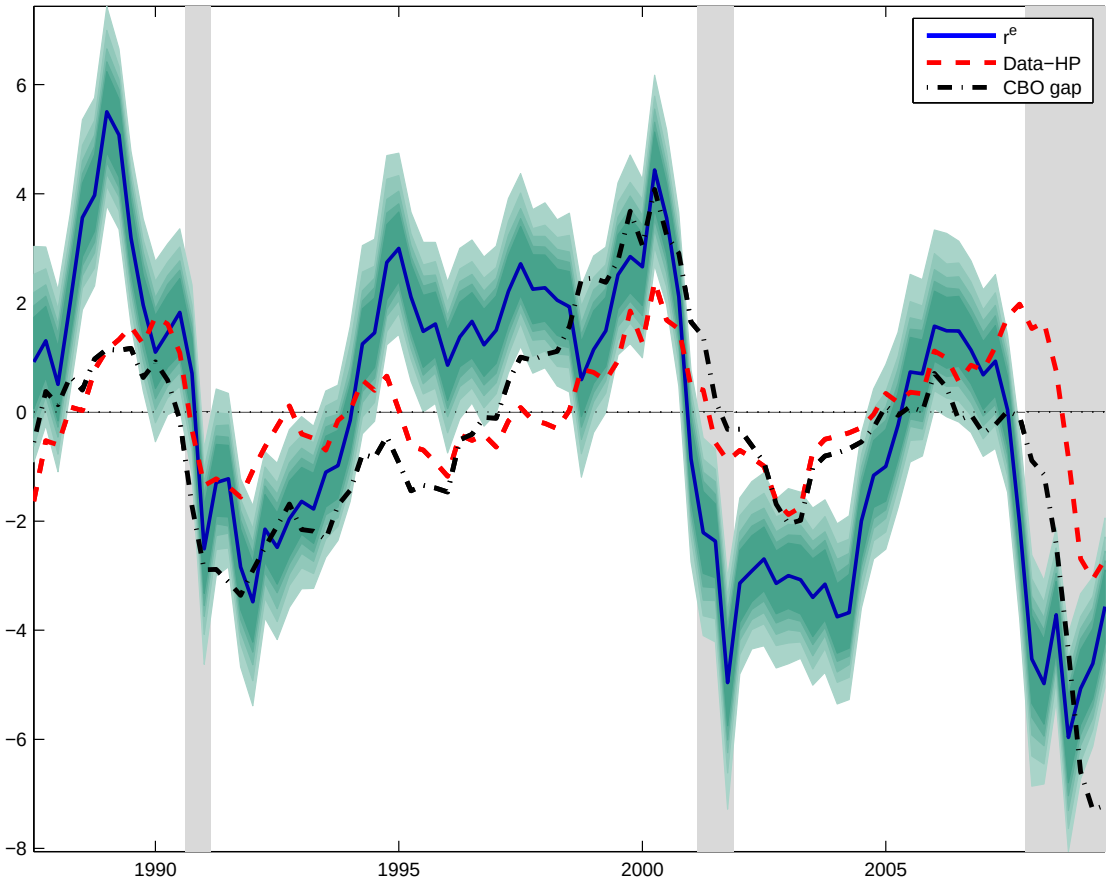
r^e
Data−HP
CBO gap
6
4
2
0
−2
−4
−6
−8
1990
1995
2000
2005

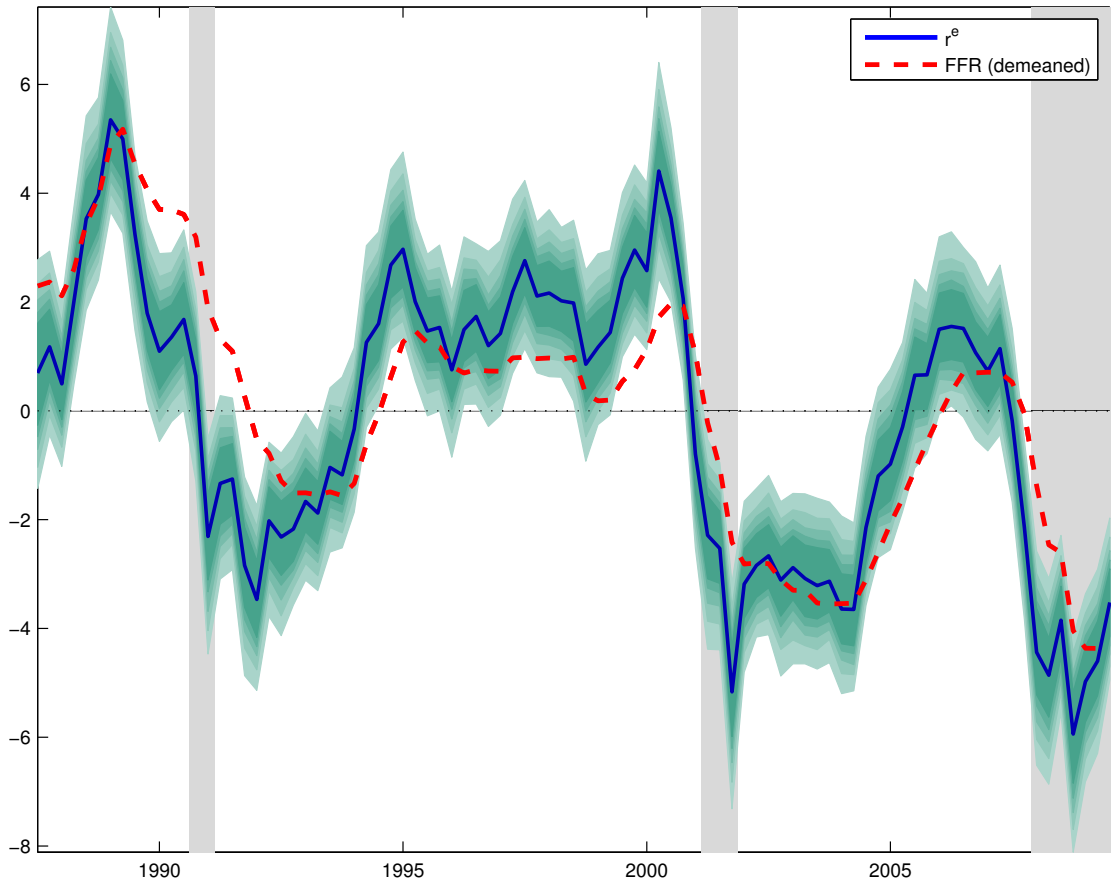

Figure 2: Evolution of model efficient real interest rate (r^e) and demeaned Federal Funds Rate (FFR demeaned), both annualized and in percentage points. The blue continuous line and the shaded area around it are the posterior median estimate of the model efficient real interest rate r^e and the 90% uncertainty bands when the interest rate rule is $i_t = \rho i_{t-1} + (1-\rho)\left(r_t^e + \phi_\pi \pi_t + \phi_x x_t^e\right) + \varepsilon_t^i$. The red dashed line is the demeaned FFR (sample mean equal to 4.5%).

Figure 3: Evolution of posterior median estimate of the model efficient real interest rate (r_t^e) across specifications Baseline ($\phi_\pi \pi_t + \phi_x x_t^e$), Re ($r_t^e + \phi_\pi \pi_t + \phi_x x_t^e$), and RePistar ($r_t^e + \pi_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_x x_t^e$).

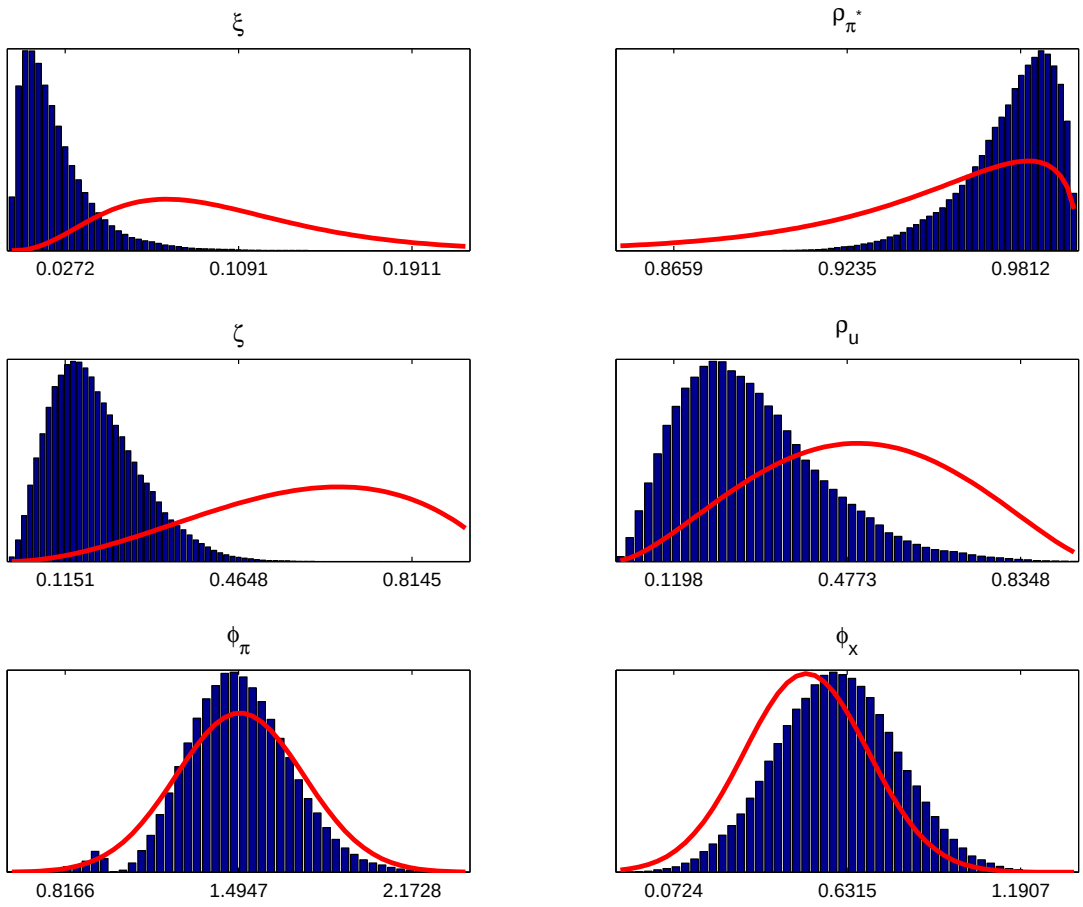

Figure 4: Prior and posterior distributions for ξ, ζ, ρ_u, ϕ_π, and ϕ_x under the RePistar specification of interest rate rule: $i_t = \rho i_{t-1} + (1-\rho)(r_t^e + \pi_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_x x_t^e) + \epsilon_t^i$. For each parameter, the solid red line represents the prior while the blue histogram is the posterior.

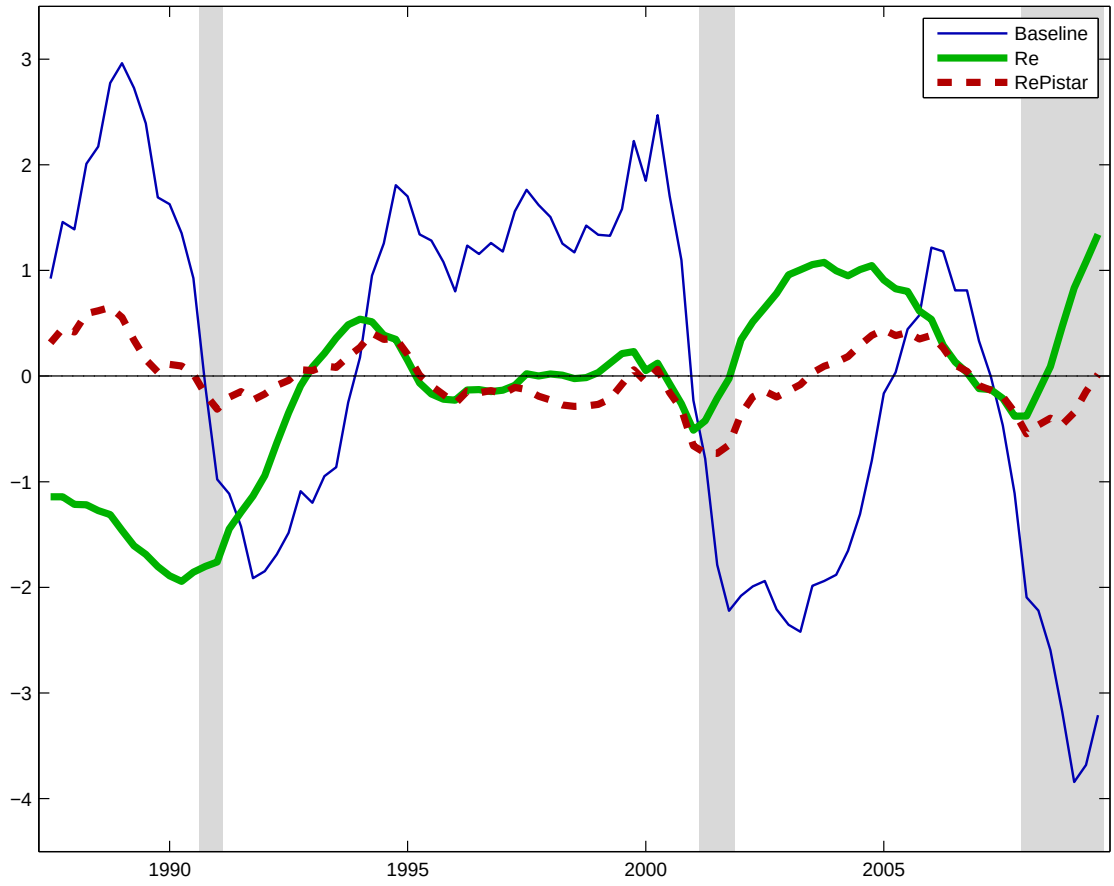

Figure 5: Evolution of the posterior median estimate of the model output gap across specifications Baseline ($\phi_\pi \pi_t + \phi_x x_t^e$), Re ($r_t^e + \phi_\pi \pi_t + \phi_x x_t^e$), and RePistar ($r_t^e + \pi_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_x x_t^e$).

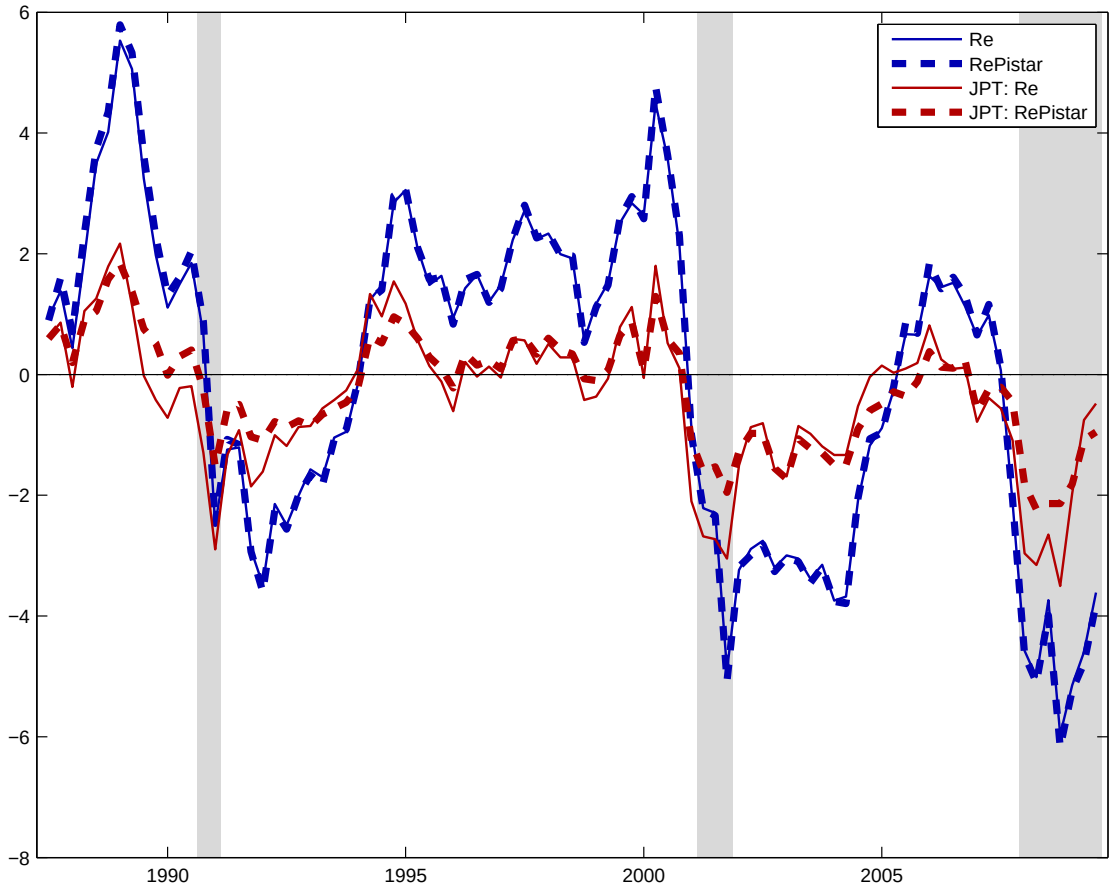

Figure 6: Evolution of posterior median estimate of the model efficient real interest rate (r_t^e) across specifications Baseline ($\phi_\pi \pi_t + \phi_x x_t^e$) and Re ($r_t^e + \phi_\pi \pi_t + \phi_x x_t^e$) for the small (dark blue) and JPT (light red) models.

Appendix A. The Model

This appendix presents the microfoundations of the model.

Appendix A.1. Households

A continuum of households of measure one populates the economy. All households, indexed by $j \in (0,1)$, discount the future at rate $\beta \in (0,1)$ and have the same instantaneous utility function, additively separable over consumption and labor, so that their objective is

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t \prod_{s=0}^{t} e^{\delta_s} \left[\log(C_t^j - \eta C_{t-1}^j) - \frac{(h_t^j)^{1+\omega}}{1+\omega} \right] \right\}.$$

The aggregate preference shock δ_t shifts the intertemporal allocation of consumption without affecting the intratemporal margin between labor and leisure.[22] We assume that δ_t follows a stationary process with mean zero of the form

$$\delta_t = \rho_\delta \delta_{t-1} + \varepsilon_t^\delta.$$

The consumption index C_t^j is a constant elasticity of substitution aggregator over differentiated goods indexed by $i \in (0,1)$

$$C_t^j \equiv \left[\int_0^1 c_t^j(i)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}. \tag{A.1}$$

Households supply their specialized labor input for the production of a specific final good. As a consequence of labor market segmentation, the wage w_t^j differs across households. However, household j can fully insure against idiosyncratic wage risk by buying at time t state-contingent securities D_{t+1}^j at price $Q_{t,t+1}$. Besides labor income, households earn after-tax Γ_t^j from ownership of the firm. The flow budget constraint for household j is

$$\int_0^1 p_t(i) c_t^j(i) di + E_t(Q_{t,t+1} D_{t+1}^j) = w_t^j h_t^j + D_t^j + \Gamma_t^j,$$

where $p_t(i)$ is the dollar price of the ith good variety.

[22] We could have also introduced a purely intratemporal shock affecting labor supply decisions only. However, in our empirical implementation of the model, hours and wages are not included among the observables. Therefore, such a shock would only affect the flexible price level of output, making it indistinguishable from a technology shock.

Appendix A.2. Firms

Firm i produces the differentiated consumption good $y_t(i)$ with a linear production function in labor

$$y_t(i) = A_t h_t(i). \tag{A.2}$$

We assume that productivity grows at rate $\gamma_t \equiv \Delta \log A_t$ and that growth rate shocks display some persistence

$$\gamma_t = \left(1 - \rho_\gamma\right)\gamma + \rho_\gamma \gamma_{t-1} + \varepsilon_t^\gamma. \tag{A.3}$$

Firms take wages as given and sell their products in monopolistically competitive goods markets, setting prices in a staggered fashion, as in Calvo (1983). Every period, independently of previous adjustments, each firm faces a probability $(1-\alpha)$ of optimally choosing its price. The α firms that do not fully optimize in a given period adjust their price according to the indexation scheme

$$p_t(i) = p_{t-1}(i)\left(\frac{P_{t-1}}{P_{t-2}}\right)^{\zeta} e^{(1-\zeta)\pi^*},$$

where P_t is the aggregate price level consistent with the consumption aggregator (A.1) and we allow for partial indexation to the long run central bank's inflation target π^*. In the event of a price change at time t, firm i chooses $p_t(i)$ to maximize the present discounted value of profits net of sales taxes τ_t

$$E_t\left\{\sum_{s=t}^{\infty} \alpha^{T-t} Q_{t,s}\left[(1-\tau_s)\, p_t(i)\left(\frac{P_{s-1}}{P_{t-1}}\right)^{\zeta} e^{(1-\zeta)\pi^*(s-t)} y_{t,s}(i) - w_s(i)\, h_s(i)\right]\right\}, \tag{A.4}$$

subject to its production function (A.2) and the demand for its own good conditional on no further price change after period t

$$y_{t,s}(i) = \left[\frac{p_t(i)}{P_s}\right]^{-\theta} Y_s, \tag{A.5}$$

where Y_t is an index of aggregate demand of the same form as (A.1).

Appendix A.3. Monetary Policy

The central bank sets the net nominal interest rate i_t with a certain degree of inertia in response to departures of aggregate demand and inflation from their respective objectives.

The non-linear formulation of the baseline interest rate rule is

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R}\right)^{\rho} \left[\left(\frac{P_t}{P_{t-1}e^{\pi^*}}\right)^{\phi_\pi} \left(\frac{Y_t}{Y_t^e}\right)^{\phi_x}\right]^{1-\rho} e^{\varepsilon_t^i}, \tag{A.6}$$

where the gross nominal interest rate is defined as

$$R_t \equiv \frac{1}{E_t Q_{t,t+1}}$$

and its average can be decomposed via the Fisher equation as $R = e^{r+\pi^*}$, which defines the steady state net real interest rate r. The continuously compounded nominal interest rate in the text is defined as $i_t \equiv \log R_t$.

Appendix B. Statistical Filters in DSGE Models

This appendix illustrates how to embed a linear filter into a dynamic rational expectation model. We begin with a brief general description of linear filtering problems. We then focus on the application to the Hodrick and Prescott (HP) filter (Hodrick and Prescott, 1997).

Appendix B.1. Linear Filters

The objective of "filtering" is to decompose the stochastic process x_t into two orthogonal components

$$x_t = y_t + \tilde{x}_t,$$

where the process y_t has power only in some frequency interval $\{(a,b) \cup (-a,-b)\} \in (-\pi,\pi)$. Then, we can represent y_t as

$$y_t = B(L)\, x_t,$$

where $B(L)$ – the *ideal* band-pass filter – is of the form

$$B(L) = \sum_{j=-\infty}^{\infty} B_j L^j.$$

Therefore, implementation of the ideal filter requires an infinite dataset. We can think about approximating the ideal filter as a projection problem. Given a sample $x = [x_1, ..., x_T]$, the estimate of $y = [y_1, ..., y_T]$ is $\hat{y} = P[y|x]$, which is of the form

$$\hat{y}_t = \sum_{=-f}^{p} \hat{B}_j^{p,f} x_{t-j},$$

where $f = T - t$ and $p = t - 1$. The main problem of this estimates is that the B coefficients require knowledge of $f_x(\omega)$, the spectral density of x.

Christiano and Fitzgerald (2003) show that, for most macro variables, the coefficients obtained by assuming that x is a random walk work quite well. One approach to the calculation of these coefficients is then to "expand" the available sample with the least squares optimal guesses of the missing data at the beginning and end of the sample. For the random walk, these data are just x_1 and x_T. Our proposal is to adopt the same philosophy (i.e. to expand the available dataset) in the context of our framework, using the rational expectations forecasts of the missing data obtained from the model.[23]

Appendix B.2. Application to the HP Filter

In this section, we discuss the application of our methodology to the HP filter. We focus on the HP filter because of its wide use in macroeconomics as a flexible device (through the choice of λ) to draw a smooth trend through the data. The HP filter provides a typical example of a "traditional" smooth measure of potential output and of the associated output gap. Its added advantage in out context is that the expression for the ideal filter is a relatively simple function of lag polynomials. The result is a parsimonious (i.e. two leads and lags) recursive representation, that requires only a modest expansion of the model's state space.

The ideal HP filter is of the form (e.g. Baxter and King, 1999)

$$\begin{aligned} HP^g &= \frac{\lambda(1-L)^2(1-F)^2}{1+\lambda(1-L)^2(1-F)^2} \\ HP^t &= \frac{1}{1+\lambda(1-L)^2(1-F)^2} \end{aligned}$$

where HP^g denotes the filter whose application results in the "gap", while HP^t denotes the filter whose application produces the trend.[24] Practical application of these filters requires an

[23] Watson (2007) proposes a similar procedure using unrestricted ARIMA processes as forecasting tools. Juillard et al. (2006) is the only example we could find of an application to DSGEs models. The main objective of all these papers is to improve the end-of-sample performance of the filters they consider.

[24] King and Rebelo (1993) originally derived these expressions as the solution of a "smoothing" problem. However, they also showed that this filter, with $\lambda = 1600$, approximates very well a high pass filter with cutoff frequency $\pi/16$ or 32 quarters.

approximation, since they embed a two-sided, infinite moving average of the data.[25] However, application of Christiano and Fitzgerald (2003) insight to a rational expectations context allows us to use the ideal filter directly, where the approximation relies on the substitution of the infinite leads and lags implicit in $HP(L)$ with rational expectation forecasts. In particular, given observations on $\log GDP_t = y_t$, we define the HP gap with parameter λ as

$$\left[1 + \lambda(1-L)^2 (1-F)^2\right] x_t^{HP(\lambda)} = \lambda(1-L)^2 (1-F)^2 y_t,$$

where now the forward and backward operators are defined by

$$\begin{aligned} Ly_t &= y_{t-1} \\ Fy_t &= E_t y_{t+1} \end{aligned}$$

as it is standard in rational expectations models (e.g. Blanchard and Fischer, 1989).

[25] Details on this approximation can be found, for example, in Baxter and King (1999).