Federal Reserve Bank of New York
Staff Reports

Fitting Observed Inflation Expectations

Marco Del Negro
Stefano Eusepi

Staff Report no. 476
November 2010

Fitting Observed Inflation Expectations
Marco Del Negro and Stefano Eusepi
Federal Reserve Bank of New York Staff Reports, no. 476
November 2010
JEL classification: C32, E21

Abstract

This paper provides evidence on the extent to which inflation expectations generated by a standard Christiano et al. (2005)/Smets and Wouters (2003)–type DSGE model are in line with what is observed in the data. We consider three variants of this model that differ in terms of the behavior of, and the public's information on, the central banks' inflation target, allegedly a key determinant of inflation expectations. We find that: 1) time-variation in the inflation target is needed to capture the evolution of expectations during the post-Volcker period; 2) the variant where agents have imperfect information is strongly rejected by the data; 3) inflation expectations appear to contain information that is not present in the other series used in estimation; and 4) none of the models fully captures the dynamics of this variable.

Key words: inflation expectations, imperfect information, Bayesian analysis, DSGE models

---

Del Negro, Eusepi: Federal Reserve Bank of New York (e-mail: marco.delnegro@ny.frb.org, stefano.eusepi@ny.frb.org). The authors thank Ippei Fujiwara for helpful comments as well as Boragan Aruoba, Giacomo Carboni, Frank Schorfheide, Chris Sims, Andrea Tambalotti, James Yetman, and seminar participants at the Board of Governors of the Federal Reserve System, De Nederlandsche Bank, the ECB, EEA 2009, the EUI, the Federal Reserve Bank of Atlanta, the Federal Reserve Bank of New York, the Riksbank, the JEDC conference on Frontiers in Structural Macroeconomic Modeling, the NBER conference on Methods and Applications for DSGE Models, SED 2009, and the UCL/CEMMAP workshop on Applied Macroeconomics and Macroeconometrics.

1 Introduction

This paper uses inflation expectations as an observable in the estimation of a DSGE model, along with a standard set of macro variables. We know very little on the extent to which DSGE models can accurately describe the behavior of observed inflation expectations.[1] The goal of the paper is therefore to provide evidence on the extent to which inflation expectations generated by standard DSGE models with nominal and real rigidities along the lines of Christiano et al. (2005), Smets and Wouters (2003), and Smets and Wouters (2007), which are currently used for policy analysis at several central banks, are in line with what observed in the data. We believe this to be an interesting question given that much of the effects of monetary policy in these models works through expectations.

We consider three variants of this prototypical DSGE model, all widely used in the literature, and we estimate them over the post-Volcker disinflation period (1984 - 2008). These variants differ in terms of the behavior of, and the agents' information on, the central banks' inflation target, allegedly a key determinant of inflation expectations. In the first variant the inflation target is fixed (as in, among others, Del Negro and Schorfheide (2009)) while in the second it is time-varying, but fully known to the public (as in Smets and Wouters (2003)).[2] We also consider a third variant where agents need to infer the time-varying target from the behavior of interest rates, as in Erceg and Levin (2003). Including this model in the analysis is a natural step, both because it is a realistic alternative to the model where agents have full information about the target (over our sample Fed officials never announced an explicit inflation target), and because Erceg and Levin (2003) argue that this type

[1] Recent literature has used survey measures of inflation expectations in the limited information estimation of New Keynesian Phillips curves (e.g. Roberts (1997), Adam and Padula (2002), and Nunes (2009)). None of this papers however studies the extent to which New Keynesian models can explain the dynamics of inflation expectations.

[2] Models where the inflation target is time-varying are ubiquitous in the estimated DSGE literature (e.g., Smets and Wouters (2003), Cogley and Sbordone (2008), and Justiniano et al. (2008)), and are also popular in macro-finance (e.g., Kozicki and Tinsley (2001), Gurkaynak et al. (2005), Rudebusch and Swanson (2008) and Dewachter (2008)).

of asymmetric information is a key feature for explaining the behavior of inflation expectations. We will refer to these three variants as the Fixed-π^*, Perfect and Imperfect Information models, respectively.

We find that a standard set of macro variables over the post-Volcker disinflation period is unable to discriminate between the Perfect and Imperfect Information models. The data slightly disfavors the Fixed-π^* version, but the evidence is not overwhelming. We also find that when we estimate the models on the dataset excluding inflation expectations and then generate a fictitious time series for expectations, for all three models the correlation between actual and model generated expectations is fairly small in levels (the median is around .25, with bands that generally include zero) and close to zero in first differences. Our baseline measure of observed inflation expectations consists of the four quarters ahead expectations for the GDP deflator obtained from the Survey of Professional Forecasters, which is the same measure used by Erceg and Levin (2003). We check for the robustness of the results using different sources of expectations and different inflation measures.

Including observed inflation expectations provides strong evidence as to which of the three models fits the data best: the Perfect Information one.[3] We show that the relative failure of the Imperfect Information model to fit observed inflation expectations is due to the fact that this model imposes much more stringent cross-equation restrictions than the Perfect Information model. Evidence based on the DSGE-VAR methodology (Del Negro and Schorfheide (2004)) confirms the above results: Whenever inflation expectations are not included the degree of misspecification (as measured by the difference in marginal likelihoods between the DSGE model and the best-fitting DSGE-VAR) is about the same across models. When this variables is included, however, the degree of misspecification for the Imperfect Information model is substantially larger than for the Perfect Information one. The DSGE-VAR evidence also suggests that even the Perfect Information model may

[3]Observed expectations are rarely formally used in the existing literature, even when comparing models that differ in the way agents form expectations. For instance, Milani (2007) compares the fit of rational expectations and learning models.

not properly capture the dynamics of observed inflation expectations: As we loosen the cross-equation restrictions the DSGE-VARs' ability to fit inflation expectations improves.

We compare the forecasting accuracy of the different models in a pseudo out-of-sample forecasting exercise. We find that the out-of-sample exercise confirms the results of Bayesian model comparison: for the dataset without expectations the Perfect and Imperfect Information models have roughly the same one period ahead forecast accuracy, while for the dataset with expectations the Perfect Information model outperforms the Imperfect Information one at that horizon. Also consistently with the model comparison findings, we show that forecasts of observed inflation expectations themselves are more accurate for the Perfect than for the Imperfect Information model. In addition we find that the four quarters ahead inflation forecasts from the Perfect Information model obtained without using observed expectations as an observable are comparable, if not marginally better, than those from the SPF. Nonetheless, adding inflation expectations to the set of observables improves the forecasting accuracy for several (but not all) variables, including inflation, real GDP growth, and interest rates, especially at longer horizons. Interestingly, this is true even for the Imperfect Information model, in spite of its documented inability to capture movements in inflation expectations.

We choose not to include the Great Disinflation period ($1981 - 1985$) in our baseline sample because of issues of structural instability: Since the early 80s the policy rule, and possibly the US economy in general, has likely changed in dimensions other than just the inflation target. Nonetheless, for completeness but also for comparison with Erceg and Levin (2003) who focus on this period, we also discuss the results including the the Great Disinflation. We find that the results for the entire sample ($1980 - 2008$) are in line with those obtained for the $1984 - 2008$ sample. Results from the Great Disinflation period only provide some weak evidence in support of the Imperfect Information model, in partial agreement with Erceg and Levin (2003).

We draw a number of conclusions from our results. First, Christiano et al.

(2005)/ Smets and Wouters (2003) -type DSGE models need time-variation in the inflation target in order to capture the evolution of expectations during the post-Volcker period to a much greater extent than they need it to fit other variables, including inflation. Second, the model where agents have imperfect information on the value of the target produces a much worse fit of inflation expectations as the model where they are fully informed. This result is somewhat surprising, as this specification was conceived precisely to explain the dynamic of inflation expectations, but can be quite intuitively explained on the ground that it imposes more stringent cross-equation restriction than the variant where agents have perfect information. These cross-equation restrictions turn out to be at odds with the data. The finding are very robust across several different specification choices, and samples.

Third, from the perspective of the econometrician inflation expectations appear to contain information that is not present in the other series. Forecasters likely have a richer information set than the econometrician using a standard set of macro variables, and including measured expectations among the observables is a way to exploit such information set.[4] This information can be exploited for both forecasting – as shown in the pseudo out-of-sample exercise – and estimating latent variables. Indeed, the result that inflation expectations generated by all the models are quite different from the actual data can be interpreted as evidence that inflation expectations bring information about latent states – such as the inflation target – which was previously unavailable to the econometrician.[5] This piece of evidence can also be interpreted as showing that this models are not well suited to fit the behavior of inflation expectations, however. The results from the DSGE-VAR approach confirm that this is certainly the case for the Imperfect Information model, but also to a lesser extent for the Perfect Information.

[4]Following the FAVAR methodology (Bernanke et al. (2005)) there are some attempts to combine factor and DSGE models with the goal of incorporating as much of the available data as possible (Boivin and Giannoni (2006), Giannone et al. (2008)). We take a different route and incorporate this information by adding agent's expectations to the list of observables.

[5]Kiley (2008) uses measures of inflation expectations, combined with an estimated monetary policy rule, to infer public beliefs about the central bank's inflation objectives.

While there are good reasons for including measured inflation expectations among the set of observables in the estimation of DSGE models (ability to discriminate across models, information content), there are also several issues associated with this choice: data revisions, timing, choice of the expectation measures. We show that the results are robust to different choices of measurement and timing assumptions, but do not directly address perhaps the most difficult issue, that of data revisions. Like recent work by Canova and Gambetti (2007), we use expectations together with revised data in our estimation, although in a robustness exercise we try to correct for the impact of revisions.

This paper focuses on a set of rational expectation models. Yet, several papers have documented that inflation expectations as measured by surveys like Michigan, Livingston and SPF fail to be consistent with rational expectations in terms of unbiasedness and serially uncorrelated forecast errors – e.g., Lloyd (1999) and Roberts (1997).[6] Why embark at all in our investigation given such evidence? Our answer is that the prominence of this class of rational expectation models both in policy making and academia, and the importance of expectation formation within these models, makes the issues addressed in this paper interesting regardless. In fact, one can see our paper as a more structural take on questions similar to those addressed in the previous literature: To what extent rational expectation models adequately explain the behavior of observed expectations, and which one comes closest? In any case, in a robustness exercise we consider deviations between model-implied and observed expectations, either using classical measurement error, or allowing for a mapping between the states of the economy and observed expectations that differs from that implied by the models.

Related to the issue of irrationality of expectations, there are several other mechanism of expectation formations, notably learning, that we do not consider here. In light of our findings, it is interesting to ask whether learning models provide a better description of observed inflation expectations than rational expectation models. In fact, very recent papers by Ormeño (2009) and Milani (2010) have tackled

[6]Some papers provide evidence in the opposite direction, however (see Rich (1989)).

this question. Both papers use survey measures of expectations to estimate a New Keynesian model with adaptive learning. In particular, Milani (2010) introduces exogenous shocks to expectations that are propagated to the economy through the learning mechanism, and shows that such shocks play an important role in explaining macroeconomic fluctuations. Ormeño (2009) shows that adaptive learning plays an important role in explaining the behavior of US inflation and inflation expectations. In Ormeño (2009) the model fit relies on the assumption that inflation expectations are formed using a small forecasting model that is not consistent with the rational expectations equilibrium. Because learning occurs using a misspecified model, the dynamics under learning and rational expectations are markedly different.

We conclude the section with a brief literature review of models where agents have imperfect information on the Central Bank's inflation target. The aforementioned study by Erceg and Levin (2003) shows that a calibrated DSGE model with asymmetric information produces a reasonable account of the Volcker disinflation and its output costs. Several other papers introduce imperfect information about the central bank's inflation target in a calibrated monetary DSGE model (Andolfatto et al. (2008), Keen (2009), and Melecky et al. (2008) among others), but do not discuss the model's ability to explain observed measures of inflation expectation. Perhaps closer to our paper, Schorfheide (2005) estimates a New Keynesian DSGE model with imperfect information using US data from 1960 to 1997. The paper finds that while the model with full information provides a better fit over the whole sample, the model with imperfect information performs better for the Volcker disinflation period. While measures of inflation expectations are not included as observables in his dataset, he compares the time series of inflation expectations generated by the two models and finds that the two models imply inflation expectations – one-year and ten-years average – that are roughly similar over the sample. Finally, Aruoba and Schorfheide (2009) study the distortionary properties of inflation and the optimal long-run inflation rate using an estimated DSGE model with a time-varying inflation target. Their dataset includes measures of long-run inflation expectations but they do not focus on the model's ability to explain inflation expectations.

The next section describes the model, with particular emphasis on the difference between perfect and imperfect information. Section 3 discusses the econometric framework for evaluating how a model estimated to fit a baseline set of time series fares in fitting an additional time series – here, inflation expectations. This is a straightforward application of Bayesian updating, which is routinely done in the DSGE estimation literature in the time series dimension, to the cross-sectional dimension. Section 4 describes the data and discusses issues related to the inclusion of measured expectations in the set of observables. Section 5 discusses our findings.

2 Model

The economy is described by a medium-scale New Keynesian model with price and wage rigidities, capital accumulation, investment adjustment costs, variable capital utilization, and habit formation. The model is based on work of Smets and Wouters (2003), Smets and Wouters (2007), and Christiano et al. (2005). The specific version is taken from Del Negro et al. (2007), except for the monetary policy rule, which we describe in detail. For brevity we only present the log-linearized equilibrium conditions and refer the reader to the above referenced papers for the derivation of these conditions from assumptions on preferences and technologies. All variables that appear subsequently are expressed as log-deviations from this steady state.

Monetary Policy: Perfect versus Imperfect Information. The central bank follows a standard feedback rule:

$$R_t = \rho_R R_{t-1} + (1 - \rho_R) \left(\psi_1(\pi_t - \pi_t^*) + \psi_2 \dot{y}_t\right) + \sigma_r \epsilon_{R,t}, \tag{1}$$

where R_t and π_t represent the interest rate and inflation, respectively, and $\dot{y}_t$ captures some measure of economic activity in log-deviations from its steady state (in the baseline specification $\dot{y}_t$ coincides with the growth rate of output), and $\epsilon_{R,t}$ is an i.i.d. shock. The inflation target π_t^*, defined in log-deviations from its non-stochastic steady state π^*, evolves according to

$$\pi_t^* = \rho_{\pi^*} \pi_{t-1}^* + \sigma_{\pi^*} \epsilon_{\pi^*,t}, \tag{2}$$

where $0 < \rho_{\pi^*} < 1$ and $\epsilon_{\pi^*,t}$ is an i.i.d. shock. We follow Erceg and Levin (2003), Smets and Wouters (2003) and Justiniano et al. (2008), among the others, and model π_t^* as following a stationary process, although our prior for ρ_{π^*} will force this process to be highly persistent. The choice of a stationary process for π_t^* is also motivated by the fact that in our sample long-term inflation expectations (10 year ahead) have moved very little in the last ten years. We view this as evidence against the random walk assumption for our post-84 sample.[7]

Under perfect information, agents observe π_t^*. Under imperfect information they need to infer the inflation target from the observed interest rate behavior (see Erceg and Levin (2003)). Call $\tilde{\pi}_t$ the residual in the feedback rule, defined as

$$\tilde{\pi}_t = (\rho_r R_{t-1} + (1-\rho_r)(\psi_1 \pi_t + \psi_2 \widehat{y}_t) - R_t)/(1-\rho_r)\psi_1. \tag{3}$$

Agents solve a signal extraction problem using

$$\tilde{\pi}_t = \pi_t^* + \sigma_T \epsilon_{R,t} \tag{4}$$

as the measurement equation (where $\sigma_T = \frac{\sigma_r}{(1-\rho_R)\psi_1}$) and (2) as the transition equation. The law of motion of $\pi^*_{t+1|t}$ is obtained using the steady state Kalman filter

$$\pi^*_{t+1|t} = \rho_{\pi^*} \pi^*_{t|t-1} + \rho_{\pi^*} K \left(\tilde{\pi}_t - \pi^*_{t|t-1} \right), \tag{5}$$

where $K = \frac{V(\frac{\sigma_{\pi^*}}{\sigma_T}, \rho_{\pi^*})}{1+V(\frac{\sigma_{\pi^*}}{\sigma_T}, \rho_{\pi^*})}$ is the steady state Kalman gain coefficient and $\sigma_T^2 V(\frac{\sigma_{\pi^*}}{\sigma_T}, \rho_{\pi^*})$ is the steady state uncertainty regarding the inflation target. V solves:

$$V = \rho_{\pi_*}^2 \left[V - V\left(V+1\right)^{-1} V \right] + (\frac{\sigma_{\pi^*}}{\sigma_T})^2.$$

We also consider the alternative law of motion for inflation target π_t^* proposed in Gurkaynak et al. (2005):[8]

$$\pi_t^* = \rho_{\pi^*} \pi^*_{t-1} + \chi \pi_{t-1} + \sigma_{\pi^*} \epsilon_{\pi^*,t}. \tag{6}$$

[7] In contrast, Cogley and Sbordone (2008) and Aruoba and Schorfheide (2009) model π_t^* as a random walk. These papers include the entire 1970s in their sample period.

[8] Ireland (2007) also consider a law of motion for the inflation target which is is affected by economic fundamentals.

As above agents know the policy rule and the evolution of the unobserved inflation target. The forecast of the unobserved inflation target $\pi^*_{t+1|t}$ (5) now becomes

$$\pi^*_{t+1|t} = \rho_{\pi_*}\pi^*_{t|t-1} + \rho_{\pi_*}K\left(\tilde{\pi}_t - \pi^*_{t|t-1}\right) + \chi\pi_t \tag{7}$$

where K is defined as before.

Finally, as an additional robustness checks we consider the case where π^*_t enters the intercept of the feedback rule, which therefore becomes:

$$R_t = \rho_R R_{t-1} + (1-\rho_R)\left(\pi^*_t + \psi_1(\pi_t - \pi^*_t) + \psi_2\dot{y}_t\right) + \sigma_r\epsilon_{R,t}. \tag{8}$$

Firms. The economy is populated by a continuum of firms that combine capital and labor to produce differentiated intermediate goods. These firms have access to the same Cobb-Douglas production function with capital elasticity α and total factor productivity Z_t. Total factor productivity is assumed to be non-stationary, and its growth rate $z_t = \ln(Z_t/Z_{t-1})$ follows the autoregressive process

$$z_t = (1-\rho_z)\gamma + \rho_z z_{t-1} + \sigma_z\epsilon_{z,t}. \tag{9}$$

Output, consumption, investment, capital, and the real wage can be detrended by Z_t. In terms of the detrended variables the model has a well-defined steady state.

The intermediate goods producers hire labor and rent capital in competitive markets and face identical real wages, w_t, and rental rates for capital, r^k_t. Cost minimization implies that all firms produce with the same capital-labor ratio

$$k_t - L_t = w_t - r^k_t \tag{10}$$

and have marginal costs

$$mc_t = (1-\alpha)w_t + \alpha r^k_t. \tag{11}$$

The intermediate goods producers sell their output to perfectly competitive final good producers, which aggregate the inputs according to a CES function. Profit maximization of the final good producers implies the following demand curve

$$y_t(j) - y_t = -\left(1 + \frac{1}{\lambda_f e^{\widetilde{\lambda}_{f,t}}}\right)(p_t(j) - p_t). \tag{12}$$

Here $y_t(j) - y_t$ and $p_t(j) - p_t$ are quantity and price for good j relative to quantity and price of the final good. The price p_t of the final good is determined from a zero-profit condition for the final good producers. We assume that the price elasticity of the intermediate goods is time-varying. Since this price elasticity affects the mark-up that intermediate goods producers can charge over marginal costs, we refer to $\widetilde{\lambda}_{f,t}$ as mark-up shock. Following Calvo (1983), we assume that in every period a fraction of the intermediate goods producers ζ_p is unable to re-optimize their prices. A fraction ι_p of these firms adjust their prices mechanically according to lagged inflation, while the remaining fraction $1 - \iota_p$ adjusts to steady state inflation π^*. All other firms choose prices to maximize the expected discounted sum of future profits, which leads to the Phillips curve:

$$\pi_t = \frac{\beta}{1 + \iota_p \beta} I\!E_t[\pi_{t+1}] + \frac{\iota_p}{1 + \iota_p \beta} \pi_{t-1} + \frac{(1 - \zeta_p \beta)(1 - \zeta_p)}{\zeta_p (1 + \iota_p \beta)} mc_t + \frac{1}{\zeta_p} \lambda_{f,t}, \tag{13}$$

where π_t is inflation and β is the discount rate. Our assumption on the behavior of firms that are unable to re-optimize their prices implies the absence of price dispersion in the steady state. As a consequence, we obtain a log-linearized aggregate production function of the form

$$y_t = (1 - \alpha) L_t + \alpha k_t. \tag{14}$$

Equations (11), (10), and (14) imply that the labor share lsh_t equals marginal costs in terms of log-deviations: $lsh_t = mc_t$.

Households. There is a continuum of households with identical preferences, which are separable in consumption, leisure, and real money balances. Households' preferences display (internal) habit formation in consumption, that is, period t utility is a function of $\ln(C_t - hC_{t-1})$, where C_t is the level of consumption. Households supply monopolistically differentiated labor services. These services are aggregated according to a CES function that leads to a demand elasticity $1 + 1/\lambda_w$ (see Equation (12)). The composite labor services are then supplied to the intermediate goods producers at real wage w_t. To introduce nominal wage rigidity, we assume that in each period a fraction ζ_w of households is unable to re-optimize their wages. A fraction ι_w of these households adjust their $t - 1$ nominal wage by $\pi_{t-1} e^{\gamma}$, where

γ represents the average growth rate of the economy, while the remaining fraction $1-\iota_p$ adjusts to steady state wage growth $\pi^* e^\gamma$. All other households re-optimize their wages. First-order conditions imply that

$$\begin{aligned} \tilde{w}_t &= \zeta_w \beta \mathbb{E}_t\Big[\tilde{w}_{t+1} + \Delta w_{t+1} + \pi_{t+1} + z_{t+1} - \iota_w \pi_t\Big] \\ &\quad + \frac{1-\zeta_w\beta}{1+\nu_l(1+\lambda_w)/\lambda_w}\left(\nu_l L_t - w_t - \xi_t + \frac{1}{1-\zeta_w\beta}\phi_t\right), \end{aligned} \tag{15}$$

where $\tilde{w}_t$ is the optimal real wage relative to the real wage for aggregate labor services, w_t, and ν_l would be the inverse Frisch labor supply elasticity in a model without wage rigidity ($\zeta_w = 0$) and differentiated labor. Moreover, ξ_t denotes the marginal utility of consumption defined below and ϕ_t is a preference shock that affects the intratemporal substitution between consumption and leisure. The real wage paid by intermediate goods producers evolves according to

$$w_t = w_{t-1} - \pi_t - z_t + \iota_w \pi_{t-1} + \frac{1-\zeta_w}{\zeta_w}\tilde{w}_t. \tag{16}$$

Households are able to insure the idiosyncratic wage adjustment shocks with state contingent claims. As a consequence they all share the same marginal utility of consumption ξ_t, which is given by the expression:

$$(e^\gamma - h\beta)(e^\gamma - h)\xi_t = -(e^{2\gamma} + \beta h^2)c_t + \beta h e^\gamma \mathbb{E}_t[c_{t+1} + z_{t+1}] + h e^\gamma (c_{t-1} - z_t), \tag{17}$$

where c_t is consumption. In addition to state-contingent claims households accumulate three types of assets: one-period nominal bonds that yield the return R_t, capital $\bar{k}_t$, and real money balances.[9]

The first order condition with respect to bond holdings delivers the standard Euler equation:

$$\xi_t = \mathbb{E}_t[\xi_{t+1}] + R_t - \mathbb{E}_t[\pi_{t+1}] - \mathbb{E}_t[z_{t+1}]. \tag{18}$$

Capital accumulates according to the law of motion

$$\bar{k}_t = (2 - e^\gamma - \delta)\big[\bar{k}_{t-1} - z_t\big] + (e^\gamma + \delta - 1)\big[i_t + S'' e^{2\gamma}(1+\beta)\mu_t\big], \tag{19}$$

[9]Since preferences for real money balances are assumed to be additively separable and monetary policy is conducted through a nominal interest rate feedback rule, money is block exogenous and we will not use the households' money demand equation in our empirical analysis.

where i_t is investment, δ is the depreciation rate of capital, and μ_t is a stochastic disturbance to the price of installed capital relative to consumption. Investment in our model is subject to adjustment costs, and S'' denotes the second derivative of the investment adjustment cost function at steady state. Optimal investment satisfies the first-order condition:

$$i_t = \frac{1}{1+\beta}\big[i_{t-1} - z_t\big] + \frac{\beta}{1+\beta}I\!E_t[i_{t+1}] + \frac{1}{(1+\beta)S''e^{2\gamma}}(\xi_t^k - \xi_t) + \mu_t, \tag{20}$$

where ξ_t^k is the value of installed capital and evolves according to:

$$\xi_t^k - \xi_t = \beta e^{-\gamma}(1-\delta)I\!E_t\big[\xi_{t+1}^k - \xi_{t+1}\big] + I\!E_t\big[(1-(1-\delta)\beta e^{-\gamma})r_{t+1}^k - (R_t - \pi_{t+1})\big]. \tag{21}$$

Capital utilization u_t is variable and r_t^k in the previous equation represents the rental rate of effective capital

$$k_t = u_t + \bar{k}_{t-1} - z_t. \tag{22}$$

The optimal degree of utilization is determined by

$$u_t = \frac{r_*^k}{a''}r_t^k. \tag{23}$$

Here a'' is the derivative of the per-unit-of-capital cost function $a(u_t)$ evaluated at the steady state utilization rate. The aggregate resource constraint is given by:

$$y_t = (1+g_*)\left[\frac{c_*}{y_*}c_t + \frac{i_*}{y_*}\left(i_t + \frac{r_*^k}{e^\gamma - 1 + \delta}u_t\right)\right] + g_t, \tag{24}$$

where c_*/y_* and i_*/y_* are the steady state consumption-output and investment-output ratios, respectively, and $g_*/(1+g_*)$ corresponds to the government share of aggregate output. The process g_t can be interpreted as exogenous government spending shock. It is assumed that fiscal policy is passive in the sense that the government uses lump-sum taxes to satisfy its period budget constraint. Finally, all stochastic processes described above are assumed to be AR(1) processes with normally distributed errors.

State-Space Representation of the DSGE Model. We use the method in Sims (2002) to solve the log-linear approximation of the DSGE model. We collect all the DSGE model parameters in the vector θ, stack the structural shocks in the vector ϵ_t, and derive a state-space representation for our vector of observables y_t, which is composed of the transition equation:

$$s_t = \mathcal{T}(\theta)s_{t-1} + \mathcal{R}(\theta)\epsilon_t, \tag{25}$$

which summarizes the evolution of the states s_t, and of the measurement equations:

$$y_t = \mathcal{Z}(\theta)s_t + \mathcal{D}(\theta), \tag{26}$$

which maps the states onto the vector of observables y_t, where $\mathcal{D}(\theta)$ represents the vector of steady state values for these observables. Specifically, for our standard set of macro time series the set of measurement equations is:

$$\begin{array}{lcl}
\text{Real output growth } {\scriptstyle (\%,\text{ annualized})} & & \\
400(\ln RGDP_t - \ln RGDP_{t-1}) & = & 400(y_t - y_{t-1} + z_t) + \gamma \\
\text{Hours } {\scriptstyle (\%)} & & \\
100 \ln L_t & = & 100(L_t + \ln L^{adj}) \\
\text{Labor Share } {\scriptstyle (\%)} & & \\
100 \ln LSH_t & = & 100(L_t + w_t - \widehat{y}_t + \ln lsh_*) \\
\text{Inflation } {\scriptstyle (\%,\text{annualized})} & & \\
400(\ln P_t - \ln P_{t-1}) & = & 400\pi_t + \pi^* \\
\text{Interest Rates } {\scriptstyle (\%,\text{annualized})} & & \\
400 \ln R_t & = & 400R_t + R_*,
\end{array} \tag{27}$$

where $RGDP_t$, L_t, LSH_t, P_t, and R_t represent real per capita GDP, total per capita hours, the labor share, the price level, and the interest rate, respectively.[10] The quantities γ, π^*, and R_* are the annualized (in percent) steady state real output growth, inflation, and nominal interest rate, respectively, lsh_* is the steady

[10]Relative to Smets and Wouters (2007) we use the labor share as opposed to growth in wages, both because it contains level information relative to wages, which among other things allows us to estimate the capital share, and because it provides a measure of real marginal costs, a key variable for inflation in New Keynesian models. We also choose to exclude consumption and investment. Del Negro et al. (2007) find that the non-stationarity of the great ratios involving consumption and investment is one of the main sources of misspecification of DSGE models. Such misspecification may in principle “pollute” the outcome of the model comparison done in this paper. Since consumption and investment as observables do not necessarily play a key role in the exercise at hand we therefore chose to exclude them. Del Negro and Schorfheide (2008), a paper about model comparison, take the same route.

states of the labor share and the parameter L^{adj} captures the units of measured hours (it can be viewed as a re-parameterization of the steady state associated with the time-varying preference parameter ϕ_t that appears in the households' utility function). As a measure of inflation we use either the GDP deflator (our baseline) or CPI, depending on the the corresponding inflation expectation measure, and as a measure of interest rates we use the federal funds rate. Appendix A provides further details on the data. In our benchmark specification we use 97 quarters of data spanning the post-Volcker disinflation period: 1984Q2 to 2008Q2.

Whenever we include observed k-quarter ahead inflation expectations $\pi_t^{O,t+k}$ to our set of time series, the set of equations (27) is augmented to include:

$$\begin{aligned} \pi_t^{O,t+k} &= \mathbb{E}_t^{\mathcal{M}_i}[\pi_{t+k}] \\ &= 400\mathcal{Z}(\theta)_{\pi,\cdot}\mathcal{T}(\theta)^k s_t + \pi^*, \end{aligned} \tag{28}$$

where $\mathbb{E}_t^{\mathcal{M}_i}[\pi_{t+k}]$ are the (annualized) inflation expectations obtained from the DSGE model $\mathcal{M}_i$. The second line shows how to compute these expectations using the transition equation (25), where $\mathcal{Z}(\theta)_{\pi,\cdot}$ is the row of $\mathcal{Z}(\theta)$ corresponding to inflation. In our application $k = 4$. Equation (28) embodies the assumption that observed expectations are rational, which arguably clashes with some of the evidence mentioned in the introduction. Section 5.8 discusses this issue, and explores alternative formulations.

3 Predictive Checks in the Cross-Section

Let $y_{1,T}^i = \{y_t^i\}_{t=1}^T$ define time series i. A natural question in the DSGE model estimation literature is the following: How does a model that is estimated to fit time series $y_{1,T}^1$ through $y_{1,T}^J$ fare in fitting time series $y_{1,T}^{J+1}$ through $y_{1,T}^{J+K}$? In this paper, for instance, we ask how the Christiano et al. (2005)/Smets and Wouters (2003) model, which allegedly fits standard macro time series well, fare in describing observed inflation expectations. The same question can be posed for asset prices, the yield curve, and several other time series.

Let $Y^0_{1,T}$ and $Y^1_{1,T}$ denote $\{y^1_{1,T},\ldots,y^J_{1,T}\}$ and $\{y^{J+1}_{1,T},\ldots,y^{J+K}_{1,T}\}$, respectively. One can of course compute the marginal likelihood for series $y^1_{1,T}$ though $y^{K+J}_{1,T}$, namely

$$p(Y^0_{1,T},Y^1_{1,T}|\mathcal{M}_i)=\int p(Y^0_{1,T},Y^1_{1,T}|\theta,\mathcal{M}_i)p(\theta|\mathcal{M}_i)d\theta, \tag{29}$$

where $\mathcal{M}_i$ is the model under consideration, $p(Y^0_{1,T},Y^1_{1,T}|\mathcal{M}_i)$ denotes the likelihood function, and $p(\theta|\mathcal{M}_i)$ the prior chosen for θ. While the quantity $p(Y^0_{1,T},Y^1_{1,T}|\theta,\mathcal{M}_i)$ is certainly of interest, and we will compute it in our application, it may not necessarily address the researcher's question. This is for two reasons. First, this quantity is often sensitive to the prior chosen, and prior elicitation for some of the DSGE model parameters can be challenging (see Del Negro and Schorfheide (2008)). The researcher who is interested in knowing how well the model fits the time series $y^{J+1}_{1,T}$ through $y^{J+K}_{1,T}$ may want to use as a prior the posterior obtained from estimating the model on time series $y^1_{1,T}$ through $y^J_{1,T}$, $p(\theta|Y^0_{1,T},\mathcal{M}_i)$, which will arguably be less dependent on the initial prior $p(\theta|\mathcal{M}_i)$ chosen. In our case, the exercise would be to use the posterior obtained from fitting standard macro time series in order to evaluate the model's ability to fit expectations. The object of interest would then be the predictive likelihood (see Geweke (2005), page 66):

$$p(Y^1_{1,T}|Y^0_{1,T},\mathcal{M}_i)=\int p(Y^1_{1,T}|\theta,Y^0_{1,T},\mathcal{M}_i)p(\theta|Y^0_{1,T},\mathcal{M}_i)d\theta, \tag{30}$$

applied to the cross sectional dimension. The second reason why we may be interested in $p(Y^1_{1,T}|Y^0_{1,T},\mathcal{M}_i)$, rather than in $p(Y^0_{1,T},Y^1_{1,T}|\mathcal{M}_i)$, is that the latter provides information on how well model $\mathcal{M}_i$ fits both $Y^0_{1,T}$ and $Y^1_{1,T}$, while the researcher may want to disentangle the goodness of fit of one set of time series versus the other. The quantity $p(Y^1_{1,T}|Y^0_{1,T},\mathcal{M}_i)$ tells us how well model $\mathcal{M}_i$ fits $Y^1_{1,T}$ only, conditional on the parameter distribution delivering the best possible fit for $Y^0_{1,T}$. This quantity easily obtains as the ratio of two objects we know how to compute, $p(Y^0_{1,T},Y^1_{1,T}|\mathcal{M}_i)$ and $p(Y^0_{1,T}|\mathcal{M}_i)$:

$$p(Y^1_{1,T}|Y^0_{1,T},\mathcal{M}_i)=\frac{p(Y^0_{1,T},Y^1_{1,T}|\mathcal{M}_i)}{p(Y^0_{1,T}|\mathcal{M}_i)}. \tag{31}$$

In addition to computing the predictive likelihood as in expression (30), it is also interesting to compare the predictive density

$$p(Y^{*1}_{1,T}|Y^0_{1,T}, \mathcal{M}_i) = \int p(Y^{*1}_{1,T}|\theta, Y^0_{1,T}, \mathcal{M}_i)p(\theta|Y^0_{1,T}, \mathcal{M}_i)d\theta \quad (32)$$

with the actual realization $Y^1_{1,T}$, where $Y^{*1}_{1,T}$ are fictitious time series for Y^1 obtained from the DSGE model $\mathcal{M}_i$ conditional on $Y^0_{1,T}$. Draws from the predictive density $p(Y^{*1}_{1,T}|Y^0_{1,T}, \mathcal{M}_i)$ are obtained using the state-space representation of the DSGE model, namely by repeatedly executing the following steps: (i) draw θ from the posterior $p(\theta|Y^0_{1,T}, \mathcal{M}_i)$, (ii) conditional on the realization of θ and $Y^0_{1,T}$, draw the unobserved states $s_{1,T}$ from $p(s_{1,T}|Y^0_{1,T}, \mathcal{M}_i)$ (e.g., using the method in Carter and Kohn (1994)), (iii) use the DSGE model-implied mapping between $s_{1,T}$ and Y^1 (in our case, expression (28)) to obtain a time series $Y^{*,1}_{1,T}$. This exercise is conducted in section 5.2. Geweke (2005) and Geweke (2007) extensively discuss the role of predictive checks in Bayesian analysis.

4 Measurement and Issues with Modeling Inflation Expectations

Several issues arise when using inflation expectations as observables in the estimation of DSGE models. First, there are many measures of inflation expectations available, for different inflation measures, and at different horizons. Our measurement choice of inflation expectations for the benchmark specification coincides with that of Erceg and Levin (2003): We use four quarters ahead expectations for the GDP deflator obtained from the Survey of Professional Forecasters. We check for the robustness of the results using different sources of expectations (Blue Chip versus SPF), and different inflation measures (CPI versus GDP deflator). An alternative source of inflation expectations is the Michigan Survey of households, which are available at the one and ten years horizons. However in that Survey households are asked about inflation in general, as opposed to any specific measure, and that makes it hard to have a measurement of expectations that is consistent with the chosen measure of inflation.

In terms of forecast horizons, we choose the longest forecast horizon for which data are available since the 1980s, namely four quarters, since arguably longer forecast horizons are more informative on agents' views about the policymakers' inflation target. Measures of inflation expectations for forecasting horizon longer than 4 quarters ahead are available but with limitations in terms of sample length and frequency. SPF provides average CPI inflation forecasts for the following 10 years but the sample starts in 1990Q4. Blue Chip and the Philadelphia Fed's Livingston survey also provide 10-years CPI inflation forecast staring 1979Q4 but the forecasts are taken only twice a year.[11]

Another serious issue is that forecasters (whether SPF or Blue Chip) have only the latest vintage of data available, while the econometrician often uses the final vintage. This is potentially important, especially for revision in the inflation measure itself, which will heavily condition the forecasts. This is not the only paper that uses inflation expectations together with revised data for macroeconomic variables (Canova and Gambetti (2007), Leduc et al. (2007), and Clark and Davig (2009), who use structural VARs, are recent examples). Addressing the issue of data revisions when using observed expectations as observables represents a major challenge, which we do not undertake in this paper. We do however show the robustness of the results when we use CPI as a measure of inflation, as opposed to the more heavily revised GDP deflator. Non-seasonally adjusted CPI is never revised. Seasonally adjusted CPI has revisions, but these are fairly small compared to those for the GDP deflator, as shown by Figure A.1 in the appendix. In addition, in a robustness exercise we try to assess whether data revisions are quantitatively important by adding the revision of time t inflation data (i.e, the difference between first and last vintage) to the measurement equation.

A third issue is the one of information synchronization. SPF forecasters pro-

[11] Concerning the 5-years horizon, Blue Chip includes forecasts which are also taken twice a year, while SPF produces quarterly forecast starting only in 2005Q3. SPF also provides quarterly 5 and 10 years forecast for PCE inflation but those start in 2007Q1. Finally, SPF produces 2 year forecasts for CPI (core and total) and PCE (core and total) inflation but they are available since 2007Q1 (CPI is available since 2005Q3).

vide their forecasts in the middle of the quarter, and hence have partial information about the state of the economy in the current quarter. We deal with this issue by checking the robustness of the results to different assumptions regarding the timing of the agents' information set. The benchmark results are obtained assuming that observed expectations are formed using current quarter information (which is also the assumption used in Canova and Gambetti (2007)). The alternative assumption, which we call "Lagged Information" specification, is that the forecasters are only endowed with information up to the previous quarter. Last, forecasts are heterogeneous, and our model cannot account for such heterogeneity (sticky information models can produce heterogeneous expectations, see Mankiw et al. (2003)). This is a very interesting and important avenue of research, which we do not pursue in this paper.

5 Results

Here is the road map of our results: The next section describes the choice of priors and their implications in terms of the observables. Section 5.2 visually investigates the extent to which inflation expectations generated from a model estimated on a standard set of macro time series resemble the actual data. Section 5.3 proceeds to a more formal model comparison exercise. Section 5.4 uses the DSGE-VAR methodology to address the issue of whether the models we consider adequately describe the behavior of inflation expectations. Section 5.5 provides the results of a pseudo out-of-sample forecasting exercise. Section 5.6 describes the features of the posterior distributions obtained under the Perfect and Imperfect Information models. Sections 5.7 and 5.8 assess the robustness of our findings. In particular, section 5.8 allows for discrepancies between model-implied and observed expectations, either due to classical measurement error, to the effect of data revisions, or to the impact of lagged forecast errors. Finally, section 5.9 considers a sample that starts in 1980 and hence includes the so-called Volcker disinflation.

5.1 Prior Choice and Prior Predictive Checks

Table 1 shows the priors for the parameters of the policy rule (1) and the associated law of motion for the inflation target π_t^* (2), which are the key parameters for the exercise conducted here. Priors for the responses to inflation (ψ_1) and the measure of economic activity (ψ_2) – output growth in the baseline specification – in the policy rule, persistence (ρ_r), and steady state inflation target (π^*) are as chosen as follows. The prior on π^* is centered using pre-sample information on inflation, as in Del Negro and Schorfheide (2008). The prior on ψ_1 and ψ_2 are centered at 2 and .2 respectively, and imply a fairly strong response to inflation and a moderate response to output. Priors on variance of i.i.d. policy shocks σ_r is centered at .15. In general the priors on the standard deviations of the shocks are chosen so that overall variance of endogenous variables is roughly close to that observed in the pre-sample 1959Q3-1984Q1, informally following the approach in Del Negro and Schorfheide (2008). Key priors are those on persistence and standard deviation of the innovation to π_t^* process, as they determine, together with the prior on σ_r, the agents' Kalman gain in the Imperfect Information model. We follow Erceg and Levin (2003) and make the process followed by π_t^* very persistent: The prior for ρ_{π^*} is centered at .95 and the 90% bands range from about .91 to .99. The half-lives of a shock to π_t^* corresponding to the prior mean and 90% bands for ρ_{π^*} are 14, 7, and 69 quarters, respectively.

In the Benchmark prior the prior on σ_{π^*}, centered at .05, is independent from all other parameters, and is fairly loose.[12] An alternative prior ("Signal-to-Noise Ratio Prior") places a prior directly on the Signal-to-Noise ratio (and hence induces dependence between σ_{π^*} and σ_r) and is centered at the value that delivers a Kalman gain of approximately .13, the value calibrated by Erceg and Levin (2003).

Priors on nominal rigidities parameters are shown in the top panel of Table 2. To check robustness to the degree of nominal rigidities in the economy we consider two priors, as in Del Negro and Schorfheide (2008): "Low Rigidities" (loosely calibrated

[12] In this and all other tables the standard deviations σ_{π^*} and σ_r are not annualized.

at Bils and Klenow (2004) values of average duration less than 2 quarters), and "High Rigidities" (duration about 4 quarters).

Priors on remaining parameters are shown in the bottom panel of Table 2. The priors on "Endogenous Propagation and Steady State" are all chosen as in Del Negro and Schorfheide (2008). Specifically, the prior for the habit persistence parameter h is centered at 0.7, which is the value used by Boldrin et al. (2001). The prior for a'' implies that in response to a 1% increase in the return to capital, utilization rates rise by 0.1 to 0.3%. These numbers are considerably smaller than the one used by Christiano et al. (2005). The 90% interval for the prior distribution on ν_l implies that the Frisch labor supply elasticity lies between 0.3 and 1.3, reflecting the micro-level estimates at the lower end, and the estimates of Kimball and Shapiro (2003) and Chang and Kim (2006) at the upper end. We use a pre-sample of observations from 1959Q3-1984Q1 to choose the prior means for the parameters that determine steady states.

The priors on standard deviations and autocorrelations are chosen so that overall variance and autocorrelations of endogenous variables is roughly close to that observed in the pre-sample 1959Q3-1984Q1 (see Table A.1 in the appendix). Table A.1 also shows that although we use the same prior for both the models under consideration – the Imperfect and Perfect Information models – the prior predictive statistics are fairly similar across models.

5.2 Inflation Expectations: Model and Data

If we generate fictional time series of expected inflation from the various models under consideration when these are estimated without including measured expectations among the observables, do we obtain anything that looks like the actual data? Figure 1 addresses this question. The figure shows the 4-quarters ahead median forecast for the GDP deflator (dashed and dotted line), together with 4-quarters ahead expected inflation generated by the Perfect Information (top panel), Imperfect Information (middle panel) models, and the model with a constant inflation target (Fixed-π^* – bottom panel). These predictive paths for inflation expectations are obtained as described in section 3. Specifically, we estimate each model $\mathcal{M}_i$ on the dataset without inflation expectations ($Y^0_{1,T}$), generate draws for the states s_t from the distribution $p(s_{1,T}|Y^0_{1,T}, \mathcal{M}_i)$, and then obtain inflation expectations using expression (28). The solid line and the shaded areas in Figure 1 represent the median, the 67, and the 95 percent bands of the predictive distribution of $I\!\!E^{\mathcal{M}_i}_t[\pi_{t+4}]$, computed period by period. In generating these figures we also condition on two lags of the variables included in Y^0, but not on the initial level of inflation expectations. In section 5.7 we show that this does not make a big difference.

Figure 1 shows that the predictive paths of inflation expectations generated by the three models are actually quite similar to one another, and quite different from the actual data. It appears that model implied inflation expectations are generally below actual expectations in the first part of the sample and above in the second part. Table 3 provides information about the correlation between model implied and actual inflation expectations. For each draw of the time series $\{I\!\!E^{\mathcal{M}_i}_t[\pi_{t+4}]\}^T_1$ we compute the correlation with $\pi^{O,t+4}_t$, and then display the median and the 90 percent bands. The correlation between model implied and actual inflation expectations is fairly small. For the whole sample the median correlation in level is around .25 for all three models, and is not significantly positive (the 5th quintile of the distribution is below zero in all three cases). In first differences the correlation is virtually zero: .12 for the Perfect Information model, and below .1 for the remaining two models.

Visually it may appear that the discrepancy between model-implied and observed expectations is largest in the first part of the sample. One may infer that the main shortcoming of the models under consideration consists in missing the timing of the downward shift in inflation expectations occurred in the early nineties. In fact, the bottom of Table 3 the correlations in the second part of the sample are even worse than for the whole sample. They are negative – and almost significantly so – in levels and close to zero in first differences.

Figure 1 is not particularly useful for the purpose of model comparison. This is because the bands in Figure 1 are computed period by period, and therefore do not account for the autocorrelation in inflation expectations. The formal model comparison exercise is therefore carried out in the next section.

5.3 Model Comparison

Table 4 shows the log marginal likelihood for three models: Imperfect Information, Perfect Information, and the model with constant inflation target (Fixed-π^*). For all models we use the Benchmark prior. The Dataset with Expectations has the SPF four quarters ahead median forecast for the GDP deflator among the observables. For these results we assume that the expectations are generated using current quarter information. In the remainder of the paper we condition on two lags of the variables included in Y^0 unless we indicate otherwise.

Table 4 shows that for the dataset without expectations (column (1)) all three models perform about the same, with the Fixed-π^* model performing slightly worse. The difference in $ln\ p(Y^0_{1,T}|\mathcal{M}_i)$ for the Imperfect and Perfect Information models is .69, which implies a posterior odd of roughly 2 in favor of the Imperfect Information model. The difference in $ln\ p(Y^0_{1,T}|\mathcal{M}_i)$ for the Fixed-π^* is larger, about 5. Although this difference implies that the posterior odds are heavily against the Fixed-π^* model, Del Negro and Schorfheide (2008) show that for marginal likelihoods for DSGE models are quite sensitive to the choice of priors, so that a difference of 5 can in principle be overturned by choosing a slightly different prior.

When SPF inflation expectations are included among the observables, the Perfect Information model performs significantly better than both the Fixed-π^* and the Imperfect Information model. The difference in the log marginal likelihoods $ln\ p(Y^0_{1,T}, Y^1_{1,T}|\mathcal{M}_i)$ between the Perfect and Imperfect Information models is about 25 in favor of the latter. The data disfavors the Fixed-π^* even more strongly. Since the marginal likelihoods $ln\ p(Y^0_{1,T}|\mathcal{M}_i)$ are similar across models, these differences translate into differences in $ln\ p(Y^1_{1,T}|Y^0_{1,T}, \mathcal{M}_i)$. They imply that the Perfect Information model fits observed inflation expectations much better than either the Imperfect Information or the Fixed-π^* model. The fact that the differences are large indicates that the extra observable included in $Y^1_{1,T}$ contains quite a lot of information as to which model describes it best.

The intuition for the above result lies in the fact that while the econometrician estimating the Perfect Information model can rely on the latent variable π^*_t to fit the data, the econometrician estimating the Imperfect Information one cannot. Recall that in the Imperfect Information model all the econometrician can infer from the data is the agents' belief about π^*_t, which we call $\pi^*_{t|t}$. The law of motion of the agents' perception of the inflation target $\pi^*_{t|t}$ is given by:

$$\pi^*_{t|t} = (1 - K)\rho_{\pi^*}\pi^*_{t-1|t-1} + K\tilde{\pi}_t, \tag{33}$$

which obtains rearranging equation (5). As we iterate this law of motion forward starting from the initial condition $\pi^*_{0|0}$, we realize that the econometricians only degree of freedom lies in the choice of this initial condition. After that, the path for $\pi^*_{t|t}$ is pinned down by that of the interest feedback rule residual $\tilde{\pi}_t$, defined in equation (3). In the benchmark specification where the interest rate responds to inflation and output growth the residual $\tilde{\pi}_t$ is pinned down by the data, for given parameters.

Figure 2 shows visually what we just described. The top panel of Figure 2 plots the mean estimate of the latent variable $\pi^*_{t|t}$ for the Imperfect Information model for the dataset without (black line) and with (gray line) inflation expectations (we condition on the same set of parameters in computing both lines, namely the value of θ that maximizes $p(\theta|Y^0_{1,T}, \mathcal{M}_i)$). Similarly, the bottom panel shows the mean estimate of the latent variable π^*_t for the Perfect Information model for the dataset without (black line) and with (gray line) inflation expectations. Since in the Imperfect Information model agents do not observe the actual π^*_t, these two latent variables are conceptually equivalent in that in each model they drive the agents' beliefs about the inflation target. Both panels also show observed inflation expectations (dashed-and-dotted line).

The time series for $\pi^*_{t|t}$ and π^*_t look very similar across the two models when the econometrician does not have information about inflation expectations (black lines in top and middle panels). Not surprisingly, for both models the movement in these time series mirrors that of the model-generated inflation expectations in Figure 1. When inflation expectations are included among the observables, the path for π_t^* in the Perfect Information model moves closer to that of observed inflation expectations. Very loosely speaking, the filtering procedure realizes that the model is failing to match the new observable, and adjusts the latent state π_t^* accordingly. For the Imperfect Information model the path for $\pi_{t|t}^*$ barely moves, and only as a result of changes in $\pi_{0|0}^*$. Because of the tight cross-equation restrictions embedded in equation (5), the filtering procedure cannot adjust $\pi_{t|t}^*$ to match inflation expectations.

5.4 DSGEs vs. VARs

How good are the DSGE models we consider at fitting observed inflation expectations? We have established in Section 5.3 that the version with Perfect Information outperforms that with Imperfect Information. Still, is the Perfect Information model any good? In order to address this question we compare the DSGE model's fit for inflation expectations to that of VARs. In particular, we follow Del Negro and Schorfheide (2004) and consider VARs with a prior that originates from the DSGE model. Del Negro and Schorfheide (2004) label this class of models DSGE-VARs (we refer to that paper for the details of this approach). Linear DSGE models such as those considered here can be viewed as – approximately – VARs subject to cross-equation restrictions. We can ask: How much does the fit of the data improve as we relax these cross-equation restrictions? The answer will give us some indication on the degree of misspecification of the DSGE model. In particular, the key parameter in the DSGE-VAR methodology is the one capturing the tightness of the DSGE model-based prior, called λ, with higher values of λ corresponding to tighter priors.

Figure 3 shows how much the fit of the VARs – as measured by the marginal likelihood – improves as we lower λ, that is, as we relax the restrictions from the DSGE model. The top and bottom panels of Figure 3 present the results for the Imperfect and Perfect Information models, respectively. Specifically, for each value of λ we compute the difference between the marginal likelihoods of the DSGE-VAR and DSGE models for the dataset with ($ln\ p_{\lambda}^{VAR}(Y_{1,T}^{0}, Y_{1,T}^{1}|\mathcal{M}_i) - ln\ p^{DSGE}(Y_{1,T}^{0}, Y_{1,T}^{1}|\mathcal{M}_i)$ – dashed line) and without expectations ($ln\ p_{\lambda}^{VAR}(Y_{1,T}^{0}|\mathcal{M}_i) - ln\ p^{DSGE}(Y_{1,T}^{0}|\mathcal{M}_i)$ – solid line). We choose a grid for the values of λ equal to $\{.33, .5, .75, 1, 1.5, 2, 5\}$, which is similar to that chosen in previous studies. We consider a VAR with two lags for this exercise.

We find that as we relax the cross-equation restrictions the marginal likelihood improves both for datasets with and without inflation expectations. This evidence, suggesting that DSGE models are misspecfied, is in line with existing results in the literature (Del Negro et al. (2007)). In all cases the λ-curve is approximately U-shaped, which is what we would expect since when the DSGE priors become "too loose" (λ approaching zero) the log marginal likelihood goes to minus infinity.

For the dataset without expectations the λ-curves (solid lines) are roughly the same: For this dataset the two models are about equally misspecified. For the dataset with expectations (dashed lines) the fit of the DSGE-VAR model relative to the DSGE model is much higher for the Imperfect Information model than for the Perfect Information, for each value of λ. In other words, when inflation expectations are included among the observables the same degree of relaxation of the cross-equation restrictions (same λ) yields much better fit for the Imperfect Information model than for the Perfect Information, indicating that the former is more at odds with the data than the latter. This finding is consistent with the results in section 5.3.

Figure 3 also shows that even the Perfect Information model is inferior to VARs in fitting the dynamics of inflation expectations. The gap between the dashed and the solid lines measures the difference between $ln\ p_{\lambda}^{VAR}(Y_{1,T}^{0}, Y_{1,T}^{1}|\mathcal{M}_i) - ln\ p_{\lambda}^{VAR}(Y_{1,T}^{0}|\mathcal{M}_i)$ and $ln\ p^{DSGE}(Y_{1,T}^{0}, Y_{1,T}^{1}|\mathcal{M}_i) - ln\ p^{DSGE}(Y_{1,T}^{0}|\mathcal{M}_i)$. The second term equals $ln\ p^{DSGE}(Y_{1,T}^{1}|Y_{1,T}^{0}, \mathcal{M}_i)$ and captures the DSGE model's ability to fit inflation expectations only, as discussed in section 3 (this number is reported in the third column of Table 4). The first term approximately measures the same quantity for the DSGE-VAR.[13] The fact that the dashed line is above the solid line, even for the Perfect Information model, and that the gap between the two lines at the peak is larger than 10 log-likelihood points, provides some evidence that even the Perfect Information model is misspecified in fitting the dynamics of inflation expectations.

5.5 Out-of-sample Forecasting Performance

Section 5.2 showed that four quarters ahead inflation expectations generated from the DSGE models, when these expectations are not part of the dataset on which the models are estimated, are quite at odds with the observed expectations. This could be due either to the fact that the DSGE models are misspecified, that is, cannot capture the dynamics of this series, or that observed expectations contain information that is absent from the other data, in the sense that they allow for a more precise reading on the state of the economy. Section 5.4 provides some evidence that the first explanation has some merit, particularly for the Imperfect Information model. In this section we perform an out-of-sample forecasting exercise in order to investigate whether the second explanation is also relevant. Namely, we want to find out whether including observed inflation expectations among the observables improves the forecasting performance for the other observables: this would suggest that observed inflation expectations are informative on the state of the economy.

Table 5 provides the root mean squared errors (RMSEs) for output, inflation, interest rates and, when part of the dataset, inflation expectations for one, four, and eight quarters ahead, computed for the period 1990Q1-2008Q2. Specifically, we estimate each model using rolling windows of 97 observations (the same number of observations for the baseline results), where the first window ends in 1990Q1 and the last one ends in 2008Q1. For each end-date we compute projections for the following eight quarters and the forecasts' mean-squared errors. Note that the four quarters ahead inflation forecasts for the dataset that include inflation coincide, by construction, with the SPF forecasts.

[13] It is not quite the same quantity because the two VARs are conditioning on a slightly different set of initial observations.

We find that for the dataset without inflation expectations the two models' RMSEs are almost identical for one period-ahead forecasts, which is consistent with the marginal likelihoods results in Table 4 (recall that the likelihood is a function of one period ahead forecast errors). Similarly, for the dataset with inflation expectations the one period-ahead RMSEs for the Perfect Information model are better than those for the Imperfect Information, particularly for output. For this dataset, the perfect Information model provides better forecasts of observed inflation expectations at all horizons, also consistently with the results in Table 4.

Interestingly, the four quarters ahead forecasts from the Perfect Information model when inflation expectations are not included are marginally more accurate than the SPF forecasts (RMSEs of 1.07 versus 1.09). Regardless, having SPF forecasts among the observables improves the forecasts of other variables, including inflation at horizons different than four quarters, for the Perfect Information model.[14] For the Imperfect Information one, this is also true with the exception of one-quarter ahead forecasts of output growth. These results suggest that observed inflation expectations provide important information to the econometrician on the state of the economy.

5.6 Posterior Estimates, Impulse Responses and Variance Decomposition

Table 6 shows the posterior mean and standard deviation (in parenthesis) of the parameters. The main differences in parameter estimates between the posterior without ($p(\theta|Y^0_{1,T}, \mathcal{M}_i)$) and with inflation expectations ($p(\theta|Y^0_{1,T}, Y^1_{1,T}, \mathcal{M}_i)$) for the Imperfect Information model are as follows. The ratio of σ_{π^*} to σ_r decreases from .13 to .11 between columns (1) ($p(\theta|Y^0_{1,T}, \mathcal{M}_i)$) and (2) ($p(\theta|Y^0_{1,T}, Y^1_{1,T}, \mathcal{M}_i)$) , and the estimates of ρ_{π^*} and ρ_r decrease as well. The importance of nominal rigidities decreases, consistently with the results discussed in section 5.7. The importance

[14] This is not true for all variables, however. Forecasting accuracy for hours and the labor share worsen at most horizons when inflation expectations are included. This is the case for both the Perfect and Imperfect Information models.

of investment adjustment increases by about 60%, which implies that investment specific shocks become much more powerful when inflation expectations are used in the estimation. The persistence of shocks all increase, except for productivity shocks where it stays about the same. The shocks standard deviations generally rise, and particularly that of government spending shocks g_t.

Changes in parameters for the Perfect Information model are quite modest. The curvature of the disutility from working ν_l decreases between columns 1 ($p(\theta|Y^0_{1,T}, \mathcal{M}_i)$) and 2 ($p(\theta|Y^0_{1,T}, Y^1_{1,T}, \mathcal{M}_i)$), thereby making hours more elastic, and the persistence of ϕ_t shock decreases (with a more elastic labor supply the reliance on ϕ_t shocks to explain movements in hours decreases). Movements in the inflation target become larger and more persistent (both σ_{π^*} and ρ_{π^*} increase).

Figure 4 shows the impulse responses of inflation and the interest rate to a one standard deviation permanent ($\epsilon_t^{\pi^*}$) and transitory (ϵ_t^R) policy shocks obtained using the posterior from the dataset without inflation expectations (impulse responses obtained from the dataset with inflation expectations are quite similar, and are shown in Figure A.2 in the Appendix). The top and bottom panels show the impulse responses under Imperfect and Perfect Information, respectively. Although the impulse responses are obtained using the respective posterior distribution, the standard deviations to both shocks are similar across models (see Table 6), so the magnitude of the shock is about the same. Yet the response of inflation and interest rates are quite different. The initial response to inflation under Imperfect Information is much smaller than under Perfect Information, as agents are uncertain as to whether it is a permanent or a transitory shock. Eventually the two impulse response converge, but this implies that the response under Imperfect Information is more hump-shaped. Conversely, the impact of a transitory shock on inflation is stronger under Imperfect Information than under Perfect Information.

Table 7 shows the (unconditional) variance decomposition computed using the posterior distribution for the Imperfect and Perfect Information models obtained using the dataset that includes observed inflation expectations. The time-varying inflation target π_t^* is the main driver of inflation expectations in the Perfect Information model, while it explains very little under Imperfect Information, consistently with the intuition discussed in section 5.3.

5.7 Robustness to the Choice of Priors, Datasets, Timing Conventions, Initial Conditions, Policy Rules, and Choice of Shocks

This section investigates the robustness of the model comparison results to the choice of priors, datasets, timing conventions, and policy rules.

Robustness to Priors: Lines (1) and (2) of Table 8 report the model comparison results under the "'High Nominal Rigidities" prior and "Signal-to-Noise Ratio" prior described in section 5.1, respectively. We find that the "High Nominal Rigidities" prior favors the Perfect Information relative to the Imperfect Information model, in that the difference in $ln\ p(Y^1_{1,T}|Y^0_{1,T}|\mathcal{M}_i)$ is larger in favor of the Perfect Information model (we use the "Low Nominal Rigidities" prior as a Benchmark precisely because it gives the Imperfect Information model the best shot). Using the "Signal-to-Noise Ratio" prior makes little difference.

Robustness to Data Sets and Timing Assumptions: Lines (3) through (6) show the log marginal likelihoods for the two models under different timing assumptions ("Lagged Information" specification), source for inflation expectations ("Blue Chip" versus SPF), and inflation measure (CPI versus GDP deflator). Under the "Lagged Information" specification the forecasters in the SPF Survey are only endowed with information up to the previous quarter. Results are robust to both timing assumptions and measurement choices. The gap in $ln\ p(Y^1_{1,T}|Y^0_{1,T}|\mathcal{M}_i)$ between the Perfect and Imperfect Information models varies among the different specifications, but is always larger than 20. The gap widens substantially whenever we use CPI (which is less subject to revisions) as opposed to the GDP Deflator.

Robustness to Conditioning Assumptions: As mentioned above, in our benchmark specification we condition on two lags of the variables included in Y^0 when computing marginal likelihoods, so that effectively we compute $ln\ p(Y^0_{1,T}|\mathcal{M}_i, Y^0_{-1,0})$

and $ln\ p(Y^0_{1,T}, Y^1_{1,T}|\mathcal{M}_i, Y^0_{-1,0})$ (for simplicity of notation we mostly omit the conditioning on $Y^0_{-1,0}$). Given that during the first part of our sample both inflation and inflation expectations are trending down, conditioning may play a non trivial role. For this reason, line (7) reports the marginal likelihoods without conditioning on any variables, while line (8) reports the results when conditioning also on the first two lags of inflation expectations. While initial conditions matter in terms of marginal likelihood computations, from the perspective of model comparison the results do not change.

Figure A.3 in the appendix is the same as Figure 1 – it shows the time series of expected inflation generated by the various models estimated without observations on actual expectations – except for the initial condition. Specifically the econometrician computing the unobserved states s_t is endowed with information about the level of inflation expectations at the beginning of the sample (1984Q1). That is, the econometrician is drawing the states s_t from the distribution $p(s_{1,T}|Y^0_{1,T}, Y^1_0, \mathcal{M}_i)$ instead of $p(s_{1,T}|Y^0_{1,T}, \mathcal{M}_i)$. For all the models the inclusion of the initial level of inflation expectations brings the model generated expectations closer to the actual data in the first few quarters of the sample, but otherwise the paths are essentially the same as those in Figure 1.

Robustness to Policy Rule Specification: Lines (9) through (12) report the model comparison results under different specifications of the policy rule, where the policy makers target the output gap (deviations from the stochastic steady state) as opposed to the output growth ("Output Gap"), a four-quarter moving average of inflation as opposed to current inflation ("4Q Inflation"), or where the the law of motion for the inflation target follows the rule suggested by Gurkaynak et al. (2005) ("GSS"). Under this rule the marginal likelihood gap between the Imperfect and Perfect Information models stays roughly constant or increases. Under the rule proposed by Gurkaynak et al. (2005) the gap narrows, but it is still larger than 17.[15] As an final robustness checks we consider the case where π^*_t enters the intercept of

[15]In the estimation of the GSS model we used the value of $\chi = .02$ in expression (6), which is the value used by Gurkaynak et al. (2005).

the feedback rule (equation (8)). Again, the results are very similar.

Robustness to Choice of Shocks (Discount Rate): The Imperfect Information models has seven shocks, as discussed in section 2. Due to the fact that two of the shocks (the i.i.d. policy shock and the shocks to the target π_t^*) are not separately observed by either the agents in the model or the econometrician, but commingle into the policy rule innovation $\tilde{\pi}_t$, effectively this model has six independent disturbances. Whenever observed inflation expectations are added to the data set, this model has therefore as many shocks as observables. Del Negro and Schorfheide (2009) show that introducing additional shocks into a model is tantamount to relaxing the cross-equation restrictions: the additional shocks can improve the model's fit by capturing dynamics that the existing set of shocks was not able to capture. One may wonder to what extent the worse fit for the imperfect information model relative to the perfect information is partly due to the set of shocks originally chosen. If so, introducing another shock may improve the imperfect information's model ability to explain inflation expectations. We therefore introduce a shock that is commonly used in DSGE models, namely a shock to the rate at which the representative agent discounts the future, which we refer to as b_t. Like most other shocks, b_t is also assumed to follow an AR(1) process. In terms of the log-linearized conditions this shock enters equation (17), which becomes:

$$\begin{aligned}(e^{\gamma}-h\beta)(e^{\gamma}-h)\xi_t &= -(e^{2\gamma}+\beta h^2)(c_t-b_t)\\ &+\beta h e^{\gamma}\mathbb{E}_t[c_{t+1}+z_{t+1}]+he^{\gamma}(c_{t-1}-z_t)-\beta h e^{-\gamma}(e^{2\gamma}+\beta h^2)\mathbb{E}_t[b_{t+1}],\end{aligned} \tag{34}$$

and equation (15), which becomes:

$$\begin{aligned}\tilde{w}_t &= \zeta_w\beta\mathbb{E}_t\Big[\tilde{w}_{t+1}+\Delta w_{t+1}+\pi_{t+1}+z_{t+1}-\iota_w\pi_{t-1}\Big]+\\ &\frac{1-\zeta_w\beta}{1+\nu_l(1+\lambda_w)/\lambda_w}\left(\nu_l L_t-w_t+\frac{e^{\gamma}(e^{\gamma}-h)}{(e^{2\gamma}+\beta h^2)}b_t-\xi_t+\frac{1}{1-\zeta_w\beta}\phi_t\right).\end{aligned} \tag{35}$$

Line (13) of Table 8 reports the marginal likelihoods under these different model specification. For the standard set of macro variables the addition of discount rate shocks barely improves the marginal likelihood, especially for the Imperfect Information model. For the data set with observed inflation expectations the improvement over the benchmark specification is substantial. This is the case for both models, however, so the relative ranking is unaffected.

Melecky et al. (2008) propose to relax in the Imperfect Information model the cross-equation restriction that links the Kalman gain in expression (5) to the signal-to-noise ratio, and hence to the estimates of the standard deviations σ_r and σ_{π^*}, and treat the Kalman gain K as an independent parameter in the estimation. We consider this modification using a prior for the K centered at .13 (the calibrated value in Erceg and Levin (2003)) with a standard deviation of .1. Line (14) of Table 8 shows that this modification has virtually no effect on the Imperfect Information model's fit, both without and with observed inflation expectation as an observable. In the latter case the posterior mean for K is .08. The intuition for this result is the same as discussed in section 5.3. Independently of how K is chosen, equation (5) embeds cross-equation restrictions that prevent any adjustment to $\pi^*_{t|t}$ required to match observed inflation expectations.

5.8 Allowing for Measurement Error/Irrationality in Observed Inflation Expectations

The measurement equation (28):

$$\pi_t^{O,t+k} = 400\mathbb{E}_t^{dsge}[\pi_{t+k}] + \pi^*$$

implies that observed inflation expectations are fully rational. In this section we ask whether our results are robust to violations of (28), whether these are due to "irrationality" of private forecasters or to issues of data revisions and data synchronization.

First, we allow for measurement error in equation (28):

$$\pi_t^{O,t+k} = 400\mathbb{E}_t^{dsge}[\pi_{t+k}] + \pi^* + \chi_t, \tag{36}$$

where the error χ_t is assumed to be either i.i.d. ("i.i.d. Meas. Error" case) or to follow and AR(1) process ("AR(1) Meas. Error" case). In both cases χ_t evolves independently from all other shocks in the model. Rows (1) and (2) of Table 9 show the marginal likelihoods for the models where we allow for measurement error in expectations.[16] The Perfect Information model is still superior to the specification with Imperfect Information when the measurement error is i.i.d.. The difference in $ln\ p(Y^1_{1,T}|Y^0_{1,T}|\mathcal{M}_i)$ is about 16, which is smaller than in Table 4 but still substantial. The fit of the two models are essentially the same under AR(1) measurement error.

We conjecture that the autoregressive measurement error largely "takes care" of the misspecification in the Imperfect (and to some extent also in the Perfect) Information model, so we essentially revert to the original result that when the dataset does not include inflation expectations the fit of the two models is about the same. We substantiate this conjecture using the variance decomposition for observed inflation expectations – both unconditional and 10-quarters ahead – shown in Table 10. We find that i.i.d. measurement error is not all that important for both the Imperfect and Perfect Information models. Its contribution is small for the unconditional variance, and between 30 and 45 percent at the 10-quarters ahead horizon. The AR(1) measurement error is the most important source of variation for observed expectations in both models, however. Measurement error explains about 60 and 40-45 percent of the variance for the Imperfect and Perfect Information models, respectively. These results may not be easily explained just by appealing to data revisions/synchronization issues. While issues of data revisions and synchronization are likely to introduce a mismatch between measured and model-generated inflation expectations, our prior would be that this mismatch is relatively short-lived. The results for the AR(1) measurement error show otherwise: For both the Perfect and Imperfect Information model the mean estimate of the AR(1) coefficient for measurement error is about .87.

Next, we investigate what lies behind the measurement error by putting more structure on the discrepancy between model implied and observed expectations.

[16]Note that the marginal likelihood for the data set without expectations $ln\ p(Y^0_{1,T}|\mathcal{M}_i)$ is the same as in the benchmark case: Whenever observed expectations are not part of the observables, the measurement error parameters do not enter the likelihood. Since the priors on these parameters are proper (they integrate to one), the marginal likelihood is the same.

That is, we use the alternative measurement equation:

$$\pi_t^{O,t+k} = 400\mathbb{E}_t^{dsge}[\pi_{t+k}] + \pi^* + \gamma' x_t, \tag{37}$$

where γ is a $\kappa \times 1$ parameter vector, and x_t consists of a $\kappa \times 1$ vector of time t observables. This alternative specification postulates that the discrepancy between model implied and observed expectations depends on observables. This is a natural case to consider for two reasons. First, the "irrationality" of observed expectations literature shows that inflation forecast errors depend on current information. Second, there is evidence that data revisions are dependent from the state of the economy (Aruoba (2008)).

The first specification we analyze is the one where x_t includes the constant: Some literature argues that a bias exists in observed inflation expectations, and the constant allows to capture such bias.[17] Row (3) of Table 9 investigates this hypothesis, using a prior for the constant has mean zero and standard deviation .75%. The marginal likelihood results indicate that the evidence in favor of a bias is very weak. The marginal likelihood is actually worse for the Imperfect Information model, and only slightly better for the Perfect Information one. The 90% posterior bands for the bias parameter are on both sides of zero.

Row (4) of Table 9 investigates whether the discrepancy depends on current inflation, output, and labor share. The prior for the elements of the vector γ all have mean zero and standard deviation .5%, and all specifications include the constant. The rationale for considering these specifications is that forecasters may react differently to the state of the economy than predicted by the rational expectations model. It turns out that none of these models significantly improves over the benchmark specification. The only parameters for which 90% posterior bands are not on both sides of zero is the response to output growth, but economically this coefficient is small.

Row (5) uses the lagged SPF forecast error (that is the difference between realized time t inflation and time $t-4$ four quarters ahead SPF forecasts) as a regressor.

[17]See for example Lloyd (1999) and Roberts (1997).

As mentioned in the introduction, several papers document that forecast errors from measured expectations display autocorrelation. We therefore check whether "correcting" observed expectations for the effect of the lagged forecast error brings the model any closer to the data. As a prior for γ we use either a fairly uninformative prior (centered at zero with a standard deviation of .5) or a prior based on the outcome of a pre-sample regression of time t on time $t-4$ SPF forecast error (this prior is centered at .5 with a standard deviation of .2). For brevity we show only the second set of results in Row (5) of Table 9. We find that the inclusion of lagged SPF forecast error in the measurement equation does not significantly improve the model's fit over the benchmark specification. The posterior mean is found to be between .1 and .2, lower than the prior mean.

Finally, in Row (6) of Table 9 we consider adding to the vector of regressors x_t the difference between the first and the last release of inflation at time t. The purpose of this exercise is to assess whether part of the discrepancy between the model's and observed expectations is due to the fact that SPF forecasters only have the latest first release of the data, while the econometrician estimating the model uses the revised data. We find that the impact of data revision is always insignificant, and that their inclusion does not improve the fit of the model.[18]

In summary, we find that the Imperfect Information model fits observed inflation expectations worse than the Perfect Information one regardless of whether we allow for a discrepancy between model implied and observed expectations. The only exception is the AR(1) measurement error. In this case the two model have roughly the same fit, but that is because the measurement error explains about half of the fluctuations in observed inflation expectations.

[18]We use an uninformative prior (centered at zero with a standard deviation of .5) for the element of γ corresponding to data revisions. This is because the outcome of the regression of SPF forecast errors on inflation revision has the opposite sign whether we use the pre-sample or the actual sample, and is barely significant whenever we include the lagged forecast error among the regressors.

5.9 Using Data from the Great Disinflation

According to Erceg and Levin (2003) the Great Disinflation of the early eighties is the poster child for the Imperfect Information model: The Central Bank raised rates in order to bring down inflation, but agents initially had trouble telling whether it represented a shift in π_t^* or a temporary interest rate shock. As a consequence, inflation expectations decline very gradually. One would therefore think that our conclusions about the ability of the Imperfect and Perfect Information model could be reversed using data from that period.

Similarly to Figure 1, Figure 5 shows the time series of expected inflation generated by the various models estimated without observations on actual expectations, except that the sample begins in 1980Q1. As in Figure 1 we see no dramatic differences between the outcomes from the Perfect and Imperfect Information models, even during the Great Disinflation period. For both models the decline in model generated expectations occurs too early relative to the decline in the actual data.

An issue with the Great Disinflation period, and particularly with its early phase, is that the rule adopted by the monetary authorities may have been different from that employed since the mid-eighties (see, among the others, Lubik and Schorfheide (2007)) . At the same time, estimating the models over the 1980-1984 period only would imply using a very short time series. For the sake of notation, call T_0, T_1, and T_2 the quarters corresponding to 1980Q1 (beginning of time series considered in Erceg and Levin), 1984Q2 (beginning of our benchmark estimation period) to 2008Q2 (end of our sample). The first row of Table 11 performs the standard model comparison exercise over the benchmark estimation period T_1 to T_2, computing the usual quantities $\ln p(Y^0_{T_1,T_2}|\mathcal{M}_i)$ and $\ln p(Y^0_{T_1,T_2}, Y^1_{T_1,T_2}|\mathcal{M}_i)$. The only difference between these numbers and those in Table 4 is that we do not condition on any presample observations (these numbers correspond to those in row (7) of Table 8). The second row of Table 11 performs the model comparison exercise over the period T_0 to T_2, and computes the quantities $\ln p(Y^0_{T_0,T_2}|\mathcal{M}_i)$ and $\ln p(Y^0_{T_0,T_2}, Y^1_{T_0,T_2}|\mathcal{M}_i)$. We find that our findings are once again fairly robust: For the standard set of observables

$Y^0_{T_0,T_2}$ the two models perform very similarly, but once observed expectations are added to the set of observables the Perfect Information model fares better.

We now ask a slightly different question. Suppose we have estimated the model over the post-84 period and made an assessment about the relative fit of the two models, how does the additional information from the Great Disinflation period update our view of the two models? The objects of interest are now

$$\ln p(Y^0_{T_0,T_1}, Y^1_{T_0,T_1} | Y^0_{T_1,T_2}, Y^1_{T_1,T_2}, \mathcal{M}_i) = \ln p(Y^0_{T_0,T_2}, Y^1_{T_0,T_2} | \mathcal{M}_i) - \ln p(Y^0_{T_1,T_2}, Y^1_{T_1,T_2} | \mathcal{M}_i),$$

where the equality shows that these objects can computed as the difference between quantities we already have computed (not conditioning on any pre-sample observations makes this convenient). Note that this exercise can be seen as a training sample *in reverse*. Usually in training sample exercises we move forward: We form a prior over the T_0 to T_1 sample and then estimate the model using data between T_1 and T_2. Here we go backward: We form a prior over the post-84 period and then compare the models over the Great Disinflation. The result is that during the Great Disinflation period the Imperfect Information model fares better, consistently with Erceg and Levin (2003) and Schorfheide (2005). The third row of Table 11 shows that the differences in log marginal likelihoods is about 10 in favor of the Imperfect Information model. This evidence should not be interpreted as indicating that the Imperfect Information model is superior to the Perfect Information one for the Great Disinflation period. It simply shows that, taken as given how bad the Imperfect Information model fits the data in the post-1984 period relative to the Perfect Information one, the 1980-84 data updates our relative assessment in favor of the Imperfect Information specification.

6 Conclusions

The paper provides evidence on the extent to which inflation expectations generated by a standard Christiano et al. (2005)/ Smets and Wouters (2003) -type DSGE model are in line with what observed in the data. We consider three variants of this model that differ in terms of the behavior of the central banks' inflation target and of the agents' information on this variable. Our findings indicate that: i) time-variation in the inflation target is needed in order to capture the evolution of expectations during the post-Volcker disinflation period; ii) the variant where agents have imperfect information is strongly rejected by the data; iii) inflation expectations appear to contain information that is not present in the other series, and iv) none of the models fully account for the evolution of observed inflation expectations.

Our findings leave several questions open. First, throughout this paper we assume that observed inflation expectations are rational. Models that incorporate adapting learning could provide a better fit to the data. Second, we only consider imperfect information about the inflation target. However, imperfect information about other driving forces can also shape the dynamics of inflation expectations. For example, many studies have discussed the link between the high US inflation rate in the 70s and the productivity slowdown that occurred at that time (and that was not fully detected until much later). Third, given our finding that observed inflation expectations contain information about the state of the economy, it may be desirable to expand the set of observables to include, for example, forecasts for GDP growth or other variables. Finally, expanding the set of macro-variables to include, for example, unemployment (and accordingly increase the complexity of the model) could improve the fit of inflation and hence of observed expectations. We leave these interesting extensions for future research.

References

Adam, Klaus and Mario Padula, "Inflation Dynamics and Subjective Expectations in the United States," CSEF Working Papers 78, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy March 2002.

Andolfatto, David, Scott Hendry, and Kevin Moran, "Are inflation expectations rational?," *Journal of Monetary Economics*, March 2008, *55* (2), 406–422.

Aruoba, S. Boragan, "Data Revisions Are Not Well Behaved," *Journal of Money, Credit and Banking*, 03 2008, *40* (2-3), 319–340.

_ **and Frank Schorfheide**, "Sticky Prices Versus Monetary Frictions: An Estimation of Policy Trade-offs," NBER Working Papers 14870, National Bureau of Economic Research, Inc April 2009.

Bernanke, Ben, Jean Boivin, and Piotr S. Eliasz, "Measuring the Effects of Monetary Policy: A Factor-augmented Vector Autoregressive (FAVAR) Approach," *The Quarterly Journal of Economics*, January 2005, *120* (1), 387–422.

Bils, Mark and Peter Klenow, "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 2004, *112*, 947–985.

Boivin, Jean and Marc Giannoni, "DSGE Models in a Data-Rich Environment," NBER Technical Working Papers 0332, National Bureau of Economic Research, Inc December 2006.

Boldrin, Michele, Lawrence Christiano, and Jonas Fisher, "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, 2001, *91*, 149–166.

Calvo, Guillermo, "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 1983, *12*, 383–398.

Canova, Fabio and Luca Gambetti, "Do expectations matter? The Great Moderation revisited," Economics Working Papers 1084, Department of Economics and Business, Universitat Pompeu Fabra November 2007.

Carter, C.K. and Robert Kohn, "On Gibbs sampling for state space models," *Biometrika*, 1994, *81* (3), 541–553.

Chang, Youngsun and Sun-Bin Kim, "Individual to Aggregate Labor Supply: A Quantitative Analysis based on Heterogeneous Agent Macroeconomy," *International Economic Review*, 2006, *47*, 1–27.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113*, 1–45.

Clark, Todd E. and Troy Davig, "Decomposing the declining volatility of long-term inflation expectations," Research Working Paper RWP 09-05, Federal Reserve Bank of Kansas City 2009.

Cogley, Timothy and Argia Sbordone, "Trend Inflation and Inflation Persistence in the New Keynesian Phillips Curve," *American Economic Review*, December 2008, *98* (5), 2101–26.

Del Negro, Marco and Frank Schorfheide, "Priors from General Equilibrium Models for VARs," *International Economic Review*, 2004, *45* (2), 643 – 673.

_ and _ , "Forming priors for DSGE models (and how it affects the assessment of nominal rigidities)," *Journal of Monetary Economics*, 2008, *55* (7), 1191 – 1208.

_ and _ , "Monetary Policy with Potentially Misspecified Models," *American Economic Review*, 2009, *99* (4), 1415–50.

_ , _ , Frank Smets, and Raphael Wouters, "On the Fit of New Keynesian Models," *Journal of Business and Economic Statistics*, 2007, *25* (2), 123 – 162.

Dewachter, Hans, "Imperfect information, macroeconomic dynamics and the yield curve : an encompassing macro-finance model," Research series 200810-19, National Bank of Belgium October 2008.

Erceg, Christopher J. and Andrew T. Levin, "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*, 2003, *50*, 915–944.

Geweke, John, *Contemporary Bayesian Econometrics and Statistics*, Wiley, 2005.

_ , "Bayesian model comparison and validation," *American Economic Review*, 2007, *97* (2), 60–64.

Giannone, Domenico, Francesca Monti, and Lucrezia Reichlin, "Incorporating Conjunctural Analysis in Structural Models," *in "The Science and Practice of Monetary Policy Today", Volker Wieland editor, Springer*, 2008.

Gurkaynak, Refet S., Brian Sack, and Eric Swanson, "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models," *American Economic Revie*, 2005, *95* (1), 425–36.

Ireland, Peter, "Changes in the Federal Reserve's Inflation Target: Causes and Consequences," *Journal of Money, Credit and Banking*, December 2007, *39* (8), 1851–82.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti, "Investment shocks and business cycles," Staff Reports 322, Federal Reserve Bank of New York 2008.

Keen, Benjamin D., "The Signal Extraction Problem Revisited: A Note on its Impact on a Model of Monetary Policy," *Macroeconomic Dynamics*, 2009, *forthcoming.*

Kiley, Michael T., "Monetary policy actions and long-run inflation expectations," Finance and Economics Discussion Series 2008-03, Board of Governors of the Federal Reserve System (U.S.) 2008.

Kimball, Miles and Matthew Shapiro, "Labor Supply: Are the Income and Substitution Effects Both Large or Both Small?," *Manuscript, University of Michigan*, 2003.

Kozicki, Sharon and P. A. Tinsley, "Shifting endpoints in the term structure of interest rates," *Journal of Monetary Economics*, June 2001, *47* (3), 613–652.

Leduc, Sylvain, Keith Sill, and Tom Stark, "Self-fulfilling expectations and the inflation of the 1970s: Evidence from the Livingston Survey," *Journal of Monetary Economics*, March 2007, *54* (2), 433–459.

Lloyd, B. Thomas Jr., "Survey Measures of Expected U.S. Inflation," *Journal of Economic Perspectives*, Fall 1999, *13* (4), 125–144.

Lubik, Thomas A. and Frank Schorfheide, "Testing for Indeterminacy: An Application to U.S. Monetary Policy: Reply," *American Economic Review*, March 2007, *97* (1), 530–533.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers, "Disagreement about Inflation Expectations," in "NBER Macroeconomics Annual 2003, Volume 18" 2003.

Melecky, Martin, Diego Rodriguez Palenzuela, and Ulf Soderstrom, "Inflation Target Transparency and the Macroeconomy," *Mimeo, Bocconi University*, 2008.

Milani, Fabio, "Expectations, Learning, and Macroeconomic Persistance," *Journal of Monetary Economics*, 2007, *54* (7), 2065 – 2082.

_ , "Expectation Shocks and Learning as Drivers of the Business Cycle," 2010. University of California, Irvine.

Nunes, Ricardo, "Inflation Dynamics: The Role of Expectations," 2009. Manuscript.

Ormeño, Arturo, "Inflation Expectations, Learning, and DSGE Model Estimation," *mimeo UPF*, 2009.

Rich, Robert W., "Testing the Rationality of Inflation Forecasts from Survey Data: Another Look at the SRC Expected Price Change Data," *The Review of Economics and Statistics*, 1989, *71* (4), 682–686.

Roberts, John M., "Is inflation sticky?," *Journal of Monetary Economics*, July 1997, *39* (2), 173–196.

Rudebusch, Glenn and Eric Swanson, "The bond premium in a DSGE model with long-run real and nominal risks," Working Paper Series 2008-31, Federal Reserve Bank of San Francisco 2008.

Schorfheide, Frank, "Learning and Monetary Policy Shifts," *Review of Economic Dynamics*, April 2005, *8* (2), 392–419.

Sims, Christopher A., "Solving Rational Expectations Models," *Computational Economics*, 2002, *20*, 1–20.

Smets, Frank and Raf Wouters, "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, 2003, *1* (5), 1123 – 1175.

_ **and** _ , "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 2007, *97* (3), 586 – 606.

A Data

The data set is obtained from Haver Analytics (Haver mnemonics are in italics). We compile observations for the variables that appear in the measurement equation (27). Real output is obtained by dividing the nominal series (*GDP*) by population 16 years and older (*LN16N*), and deflating using the chained-price GDP deflator (*JGDP*). We compute quarter-to-quarter output growth as log difference of real GDP per capita and multiply the growth rates by 400 to convert them into annualized percentages. Our measure of hours worked is computed by taking total hours worked reported in the National Income and Product Accounts (NIPA), which is at annual frequency, and interpolating it using growth rates computed from hours of all persons in the non-farm business sector (*LXNFH*). We divide hours worked by *LN16N* to convert them into per capita terms. We then take the log of the series multiplied by 100 so that all figures can be interpreted as percentage changes in hours worked. The labor share is computed by dividing total compensation of employees (*YCOMP*) obtained from the NIPA by nominal GDP. We then take the log of the labor share multiplied by 100. Inflation rates are defined as log differences of the GDP deflator and converted into annualized percentages. The nominal rate corresponds to the effective Federal Funds Rate (*FFED*), also in percent. As an alternative measure of the nominal rate we use the three months Tbill (*FTBS3*),

We use Survey of Professional Forecasters (SPF) quarterly measures of expected inflation. We consider both expectations for GDP deflator[19] and for CPI inflation. In particular, we use the median four -quarters-ahead forecast of inflation in annualized terms. Concerning the information available to the forecasters, the survey is sent out at the end of the first month of each quarter and responses deadlines occur in the middle month of each quarter. Therefore, respondents have knowledge about the BEA advance report of the National Income and Product Accounts. We also compute the revisions in GDP deflator and CPI occurred since 1982 using the real time dataset available from the Federal Reserve Bank of Philadelphia

[19]In more detail, the forecast are for the GDP price index, seasonally adjusted (base year varies). Prior to 1996, the forecast variable was the GDP implicit deflator. Prior to 1992, the GNP deflator.

(http://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/).

As an alternative measure of inflation expectations, we use Blue Chip monthly forecasts of CPI inflation. We choose forecast horizons of three and four quarters ahead. In order to compare Blue Chip and SPF quarterly forecast of CPI inflation, we use the Blue Chip forecasts available in the middle month of each quarter. This roughly corresponds to the time period when SPF participants provide their forecasts.

Table 1: Priors on Policy Parameters

Parameter	Domain	Density	Para (1)	Para (2)	5%	95%
ψ_1	$\mathbb{R}^+$	Gamma	2.00	0.25	1.592	2.408
ψ_2	$\mathbb{R}^+$	Gamma	0.20	0.10	0.049	0.349
ρ_r	[0,1)	Beta	0.50	0.200	0.170	0.827
π^*	$\mathbb{R}$	Normal	4.3	2.5	0.520	8.17
σ_r	$\mathbb{R}^+$	InvGamma	0.150	4.00	0.080	0.298
ρ_{π^*}	[0,1)	Beta	0.950	0.025	0.913	0.989
Benchmark Prior						
σ_{π^*}	$\mathbb{R}^+$	InvGamma	0.050	8.000	0.032	0.078
Signal-to-Noise Ratio Prior						
$\sigma_{NR} = \frac{\sigma_P}{\sigma_T}$	$\mathbb{R}^+$	Gamma	0.180	0.150	0.001	0.380

Notes: Para (1) and Para (2) correspond to means and standard deviations for the Beta, Gamma, and Normal distributions and to s and ν for the Inverse Gamma distribution, where $p_{\mathcal{IG}}(\sigma|\nu, s) \propto \sigma^{-\nu-1}e^{-\nu s^2/2\sigma^2}$. The last two columns report the 5^{th} and 95^{th} quintile of the prior distribution.

Table 2: Priors on Non-Policy Parameters

Parameter	Domain	Density	Para (1)	Para (2)	5%	95%
Priors on Nominal Rigidities Parameters						
Low Rigidities (Benchmark)						
ζ_p	[0,1)	Beta	0.450	0.100	0.285	0.614
ζ_w	[0,1)	Beta	0.450	0.100	0.285	0.614
High Rigidities						
ζ_p	[0,1)	Beta	0.750	0.100	0.590	0.913
ζ_w	[0,1)	Beta	0.750	0.100	0.590	0.913
Priors on "Endogenous Propagation and Steady State" Parameters						
α	[0,1)	Beta	0.330	0.020	0.297	0.362
s''	$\mathbb{R}^+$	Gamma	4	1.500	1.614	6.303
h	[0,1)	Beta	0.700	0.050	0.619	0.782
a'	$\mathbb{R}^+$	Gamma	0.200	0.100	0.049	0.349
ν_l	$\mathbb{R}^+$	Gamma	2	0.75	0.787	3.137
r^*	$\mathbb{R}^+$	Gamma	1.5	1	0.106	2.883
γ	$\mathbb{R}^+$	Gamma	1.650	1	0.204	3.073
g^*	$\mathbb{R}^+$	Gamma	0.300	0.100	0.143	0.459
ι_p	[0,1)	Beta	0.5	0.280	0.043	0.922
ι_w	[0,1)	Beta	0.5	0.280	0.049	0.932
Priors on ρs and σs						
ρ_z	[0,1)	Beta	0.400	0.250	0.000	0.764
ρ_ϕ	[0,1)	Beta	0.750	0.150	0.530	0.982
ρ_{λ_f}	[0,1)	Beta	0.750	0.150	0.530	0.982
ρ_μ	[0,1)	Beta	0.750	0.150	0.530	0.982
ρ_g	[0,1)	Beta	0.750	0.150	0.530	0.982
σ_z	$\mathbb{R}^+$	InvGamma	0.200	4.000	0.107	0.395
σ_ϕ	$\mathbb{R}^+$	InvGamma	2.500	4.000	1.326	4.930
σ_{λ_f}	$\mathbb{R}^+$	InvGamma	0.300	4.000	0.161	0.596
σ_μ	$\mathbb{R}^+$	InvGamma	0.500	4.000	0.264	0.99
σ_g	$\mathbb{R}^+$	InvGamma	0.300	4.000	0.159	0.594

Notes: Para (1) and Para (2) correspond to means and standard deviations for the Beta, Gamma, and Normal distributions and to s and ν for the Inverse Gamma distribution, where $p_{\mathcal{IG}}(\sigma|\nu,s) \propto \sigma^{-\nu-1}e^{-\nu s^2/2\sigma^2}$. The last two columns report the 5^{th} and 95^{th} quintile of the prior distribution.

Table 3: Correlation between Model-Implied and SPF Inflation Expectation

	Perfect Information		Imperfect Information		Fixed π^*	
Entire Sample (1984Q2-2008Q2)						
Level	0.25	(-0.06 , 0.52)	0.27	(-0.03 , 0.54)	0.22	(-0.27 , 0.60)
First Difference	0.12	(-0.02 , 0.25)	0.07	(-0.08 , 0.20)	0.08	(-0.05 , 0.21)
Second Half of Sample (1996Q1-2008Q2)						
Level	-0.17	(-0.41 , 0.07)	-0.28	(-0.54 , 0.06)	-0.28	(-0.54 , 0.06)
First Difference	0.06	(-0.12 , 0.24)	0.06	(-0.13 , 0.24)	0.07	(-0.08 , 0.19)

Notes: The table reports the correlation between predictive paths for four quarters ahead inflation expectations, obtained as described in section 3, and SPF inflation expectations. The former are computed using the dataset that does not include inflation expectations among the observables. We report the median and in parenthesis the 10^{th} and 90^{th} quintile of the posterior distribution.

Table 4: Model Comparison

	$\ln p(Y^0_{1,T})$	$\ln p(Y^0_{1,T}, Y^1_{1,T})$	$\ln p(Y^1_{1,T} \vert Y^0_{1,T})$
	Dataset without Expectations	Dataset with Expectations	
	(1)	(2)	(2) - (1)
Imperfect Information	-703.62	-811.04	-107.42
Perfect Information	-704.31	-786.35	-82.04
Fixed π^*	-709.29	-821.84	-112.55

Notes: The table shows the log marginal likelihood for three models: Imperfect Information, Perfect Information, and the model with constant inflation target (Fixed-π^*). For all models we use the Benchmark prior. The Dataset with Expectations uses the SPF four quarters ahead median forecast for the GDP deflator. We assume that the expectations are generated using current quarter information.

Table 5: Out-of-Sample Forecasting Performance

Quarters ahead	Output Growth	Inflation	Interest Rate	Exp. Inflation
	Imperfect Information			
	w/o Infl. Exp			
1	2.67	1.02	1.73	
4	2.44	1.24	2.86	
8	2.56	1.79	3.45	
	with Infl. Exp			
1	3.09	0.91	1.32	0.45
4	2.01	1.09	1.95	0.81
8	2.22	1.52	2.31	1.19
	Perfect Information			
	w/o Infl. Exp			
1	2.67	1.04	1.79	
4	2.36	1.07	2.84	
8	2.52	1.57	3.19	
	with Infl. Exp			
1	2.60	0.91	1.31	0.44
4	2.18	1.09	2.07	0.71
8	2.41	1.43	2.60	0.94

Notes: The table provides the root mean squared errors (RMSEs) for output, inflation, interest rates and, when part of the dataset, inflation expectations for 1, 4, and 8 quarters ahead, computed for the period 1990Q1-2008Q2.

Table 6: Posterior Estimates for Selected Parameters

Parameters	**Imperfect Information** Dataset without Expectations (1)	**Imperfect Information** Dataset with Expectations (2)	**Perfect Information** Dataset without Expectations (3)	**Perfect Information** Dataset with Expectations (4)
Policy Parameters				
ψ_1	2.442 (0.225)	1.915 (0.123)	2.497 (0.247)	2.324 (0.191)
ψ_2	0.282 (0.112)	0.255 (0.106)	0.232 (0.093)	0.264 (0.110)
ρ_r	0.407 (0.077)	0.375 (0.065)	0.454 (0.067)	0.592 (0.043)
$\rho_{\pi^\star}$	0.945 (0.025)	0.907 (0.021)	0.943 (0.025)	0.974 (0.011)
σ_r	0.404 (0.037)	0.422 (0.033)	0.389 (0.036)	0.435 (0.035)
σ_{π^*}	0.054 (0.010)	0.048 (0.009)	0.058 (0.012)	0.066 (0.009)
Nominal Rigidities Parameters				
ζ_p	0.579 (0.061)	0.530 (0.057)	0.558 (0.051)	0.580 (0.061)
ι_p	0.285 (0.182)	0.494 (0.202)	0.346 (0.181)	0.317 (0.167)
ζ_w	0.249 (0.069)	0.186 (0.031)	0.238 (0.061)	0.353 (0.098)
ι_w	0.400 (0.251)	0.540 (0.257)	0.375 (0.253)	0.370 (0.236)
Other "Endogenous Propagation and Steady State" Parameters				
α	0.340 (0.003)	0.340 (0.004)	0.340 (0.003)	0.341 (0.003)
s''	2.831 (0.880)	4.529 (1.152)	3.002 (0.902)	3.543 (1.205)
h	0.649 (0.047)	0.636 (0.053)	0.658 (0.049)	0.640 (0.046)
a''	0.291 (0.112)	0.212 (0.097)	0.275 (0.102)	0.274 (0.095)
ν_l	2.153 (0.534)	2.690 (0.649)	2.271 (0.588)	1.327 (0.510)
r^*	1.000 (0.423)	1.424 (0.541)	1.019 (0.452)	1.259 (0.471)
π^*	2.470 (0.996)	3.068 (0.574)	2.106 (0.759)	3.662 (1.134)
γ	1.629 (0.333)	1.511 (0.330)	1.646 (0.362)	1.454 (0.314)
g^*	0.272 (0.090)	0.304 (0.100)	0.287 (0.092)	0.306 (0.107)
ρs and σs				
ρ_z	0.203 (0.094)	0.200 (0.095)	0.247 (0.090)	0.177 (0.098)
ρ_ϕ	0.837 (0.071)	0.980 (0.013)	0.850 (0.062)	0.569 (0.218)
ρ_{λ_f}	0.823 (0.073)	0.838 (0.059)	0.840 (0.058)	0.803 (0.071)
ρ_μ	0.885 (0.050)	0.910 (0.025)	0.897 (0.044)	0.894 (0.051)
ρ_g	0.810 (0.116)	0.824 (0.056)	0.798 (0.140)	0.982 (0.016)
σ_z	0.699 (0.055)	0.693 (0.052)	0.709 (0.055)	0.689 (0.047)
σ_ϕ	3.008 (0.516)	3.327 (0.589)	3.055 (0.638)	2.656 (0.660)
σ_{λ_f}	0.146 (0.031)	0.175 (0.031)	0.156 (0.026)	0.149 (0.023)
σ_μ	0.468 (0.115)	0.410 (0.083)	0.464 (0.111)	0.398 (0.099)
σ_g	0.291 (0.050)	0.426 (0.047)	0.267 (0.050)	0.410 (0.050)

Notes: The table reports the posterior mean and standard deviation (in parenthesis) of the parameters for the Imperfect and Perfect Information models obtained from both the datasets with and without inflation expectations.

Table 7: Variance Decomposition

Variables	Tech	ϕ	μ	g	λ_f	π^*	Money
Imperfect Information							
Output Growth	0.25	0.35	0.11	0.22	0.05	0.00	0.01
Labor Supply	0.00	0.94	0.05	0.01	0.01	0.00	0.00
Labor Share	0.05	0.03	0.00	0.02	0.88	0.00	0.01
Inflation	0.13	0.15	0.44	0.07	0.08	0.03	0.07
Interest Rate	0.08	0.09	0.60	0.08	0.05	0.00	0.08
Exp. Inflation	0.01	0.01	0.90	0.00	0.00	0.05	0.00
Perfect Information							
Output Growth	0.29	0.12	0.17	0.20	0.10	0.00	0.04
Labor Supply	0.03	0.09	0.31	0.3	0.06	0.00	0.01
Labor Share	0.06	0.07	0.00	0.01	0.83	0.00	0.02
Inflation	0.06	0.08	0.11	0.01	0.08	0.59	0.05
Interest Rate	0.05	0.08	0.35	0.01	0.06	0.28	0.14
Exp. Inflation	0.01	0.00	0.14	0.00	0.00	0.84	0.00

Notes: The table shows the (unconditional) variance decomposition computed using the posterior distribution for the Imperfect and Perfect Information models obtained using the dataset that includes observed inflation expectations.

Table 8: Robustness of Model Comparison Results

Imperfect Information			Perfect Information		
$\ln p(Y^0_{1,T})$	$\ln p(Y^0_{1,T}, Y^1_{1,T})$	$\ln p(Y^1_{1,T} \mid Y^0_{1,T})$	$\ln p(Y^0_{1,T})$	$\ln p(Y^0_{1,T}, Y^1_{1,T})$	$\ln p(Y^1_{1,T} \mid Y^0_{1,T})$
Dataset without Expectations	Dataset with Expectations		Dataset without Expectations	Dataset with Expectations	
(1)	(2)	(2) - (1)	(3)	(4)	(4) - (3)
Robustness to Priors					
(1) High Nominal Rigidities Prior					
-701.65	-820.84	-119.19	-705.39	-789.26	-83.87
(2) Signal-to-Noise Ratio Prior					
-703.86	-811.97	-108.11	-709.66	-786.59	-76.93
Robustness to Data Sets and Timing Assumptions					
(3) Lagged Information					
-703.62	-800.74	-97.12	-704.31	-780.53	-76.22
(4) Blue Chip Expectations					
-703.62	-761.68	-58.06	-704.31	-742.11	-37.80
(5) CPI and SPF Expectations					
-761.28	-844.98	-83.70	-763.72	-771.38	-7.66
(6) CPI and Blue Chip Expectations					
-761.28	-865.04	-103.76	-763.72	-779.31	-15.59
Robustness to Conditioning Assumptions					
(7) No Conditioning					
-711.641	-816.67	-105.03	-711.67	-789.84	-78.17
(8) Conditioning on Initial Level of Inflation Expectations					
	-810.49			-784.83	

Notes: The table shows the log marginal likelihood for the Imperfect Information and Perfect Information models under different choices of priors, datasets, timing conventions, policy rules, and set of shocks. Lines (1) and (2) report the results under the "High Nominal Rigidities" prior and "Signal-to-Noise Ratio" prior, respectively. Lines (3) to (6) show the log marginal likelihood for the two models under different timing assumptions ("Lagged Information" specification), measures of inflation and measures of inflation expectations ("Blue Chip Expectations", "CPI and SPF Expectations", "CPI and Blue Chip Expectations"). Lines (7) and (8) report the results under different conditioning assumptions. Lines (9)-(11) report the results under different specifications of the policy rule, where the policy makers target output growth as opposed to the output gap ("Output Growth"), a four-quarter moving average of inflation as opposed to current inflation ("4Q Inflation"), or where the the law of motion for the inflation target follows the rule suggested by Gurkaynak et al. (2005) ("GSS"). Line (12) consider the case where π^*_t enters the intercept of the feedback rule, as in equation (8). Line (13) reports the results after augmenting the model with an additional shock (discount rate shock). Line (14) uses the specification proposed by Melecky et al. (2008)

Table 8: Robustness of Model Comparison Results – Continued

Imperfect Information			Perfect Information		
$\ln p(Y^0_{1,T})$	$\ln p(Y^0_{1,T}, Y^1_{1,T})$	$\ln p(Y^1_{1,T} \mid Y^0_{1,T})$	$\ln p(Y^0_{1,T})$	$\ln p(Y^0_{1,T}, Y^1_{1,T})$	$\ln p(Y^1_{1,T} \mid Y^0_{1,T})$
Dataset without Expectations	Dataset with Expectations		Dataset without Expectations	Dataset with Expectations	
(1)	(2)	(2) - (1)	(3)	(4)	(4) - (3)
Robustness to Policy Rule Specification					
(9) Output Gap					
-715.46	-816.23	-100.77	-709.17	-791.74	-82.57
(10) 4Q Inflation					
-703.74	-820.96	-117.22	-698.88	-790.42	-91.5
(11) GSS					
-707.79	-805.64	-97.85	-709.45	-789.99	-80.54
(12) π^*_t entering intercept					
-702.78	-806.99	-104.21	-704.99	-787.41	-82.42
Other Robustness Checks					
(13) Additional Shocks (Discount Rate)					
-701.09	-792.10	-90.01	-703.58	-773.04	-69.46
(14) Melecky et al. (2008)					
-704.59	-811.68	-107.09			

Table 9: Allowing for Measurement Error/Irrationality in Observed Inflation Expectations

Imperfect Information			Perfect Information		
$\ln p(Y^0_{1,T})$	$\ln p(Y^0_{1,T}, Y^1_{1,T})$	$\ln p(Y^1_{1,T} \| Y^0_{1,T})$	$\ln p(Y^0_{1,T})$	$\ln p(Y^0_{1,T}, Y^1_{1,T})$	$\ln p(Y^1_{1,T} \| Y^0_{1,T})$
Dataset without Expectations	Dataset with Expectations		Dataset without Expectations	Dataset with Expectations	
(1)	(2)	(2) - (1)	(3)	(4)	(4) - (3)
(0) Benchmark specification					
-703.62	-811.04	-107.42	-704.31	-786.35	- 82.04
Measurement Error					
(1) i.i.d. Measurement Error					
-703.62	-796.31	-92.69	-704.31	-780.89	-76.58
(2) AR(1) Measurement Error					
-703.62	-775.31	-71.69	-704.31	-775.21	-70.90
(3) Bias					
-703.62	-811.48	-107.86	-704.31	-784.82	-80.51
(4) Bias + Response to Current Inflation, Labor Share, and Output Growth					
-703.62	-814.29	-110.67	-704.31	-786.16	-81.85
(5) Response to Lagged Forecast Error					
-703.62	-812.60	-108.98	-704.31	-786.22	-81.91
(6) Response to Lagged Forecast Error and Data Revision					
-703.62	-815.23	-111.61	-704.31	-789.55	-85.24

Notes: The table shows the log marginal likelihood for the Imperfect Information and Perfect Information models when allowing for discrepancies between observed and model generated expectations. Line (0) shows the results from the benchmark specification for ease of comparison. Lines (1) and (2) report the log marginal likelihood for the two models measurement errors are added ("i.i.d. Measurement Error", and "AR(1) Measurement Error"). Lines (3) and (4) allow for bias in measured inflation expectations, and for a different response to current inflation, labor share, and output growth than that warranted by the rational expectation model. Lines (5) and (6) introduce in the measurement equation for observed inflation expectation a correction to lagged forecast error and data revision.

Table 10: Variance Decomposition for Observed Inflation Expectations: Models with Measurement Errors

Variables	Tech	ϕ	μ	g	λ_f	π^*	meas.	Money
Unconditional								
Imperfect Information								
i.i.d. Meas. Error	0.02	0.02	0.63	0.00	0.01	0.16	0.14	0.01
AR(1) Meas. Error	0.01	0.01	0.26	0.00	0.01	0.12	0.57	0.00
Perfect Information								
i.i.d. Meas. Error	0.01	0.01	0.21	0.00	0.01	0.67	0.07	0.00
AR(1) Meas. Error	0.01	0.00	0.25	0.00	0.00	0.27	0.41	0.00
Ten Quarters Ahead								
Imperfect Information								
i.i.d. Meas. Error	0.01	0.01	0.60	0.00	0.01	0.16	0.18	0.01
AR(1) Meas. Error	0.01	0.01	0.27	0.00	0.01	0.09	0.60	0.01
Perfect Information								
i.i.d. Meas. Error	0.01	0.02	0.29	0.00	0.02	0.53	0.11	0.00
AR(1) Meas. Error	0.01	0.00	0.28	0.00	0.00	0.24	0.43	0.00

Notes: The table shows the posterior means of the variance decomposition for observed inflation expectations – both unconditional and ten quarters ahead – for the Imperfect Information and Perfect Information models with both i.i.d. and AR(1) measurement error. The posteriors are obtained using the dataset that includes observed inflation expectations.

Table 11: Using Data from the Great Disinflation

Imperfect Information			Perfect Information		
$\ln p(Y^0_{T_i,T_j})$	$\ln p(Y^0_{T_i,T_j}, Y^1_{T_i,T_j})$	$\ln p(Y^1_{T_i,T_j} \mid Y^0_{T_i,T_j})$	$\ln p(Y^0_{T_i,T_j})$	$\ln p(Y^0_{T_i,T_j}, Y^1_{T_i,T_j})$	$\ln p(Y^1_{T_i,T_j} \mid Y^0_{T_i,T_j})$
Dataset without Expectations	Dataset with Expectations		Dataset without Expectations	Dataset with Expectations	
(1)	(2)	(2) - (1)	(3)	(4)	(4) - (3)
(1) Post-Disinflation Data Set (1984-2008, standard)					
-711.641	-816.67	-105.03	-711.67	-789.84	-78.17
(2) Post-1980 Data Set (1980-2008)					
-918.02	-1057.45	-139.43	-921.38	-1041.51	-120.13
(3): (2)-(1), Updating over the Great Disinflation Period					
	-240.78			-251.67	

Notes: The table shows the log marginal likelihood for the Imperfect Information and Perfect Information models under different data sets.

Figure 1: Inflation Expectations: Data vs Model Prediction

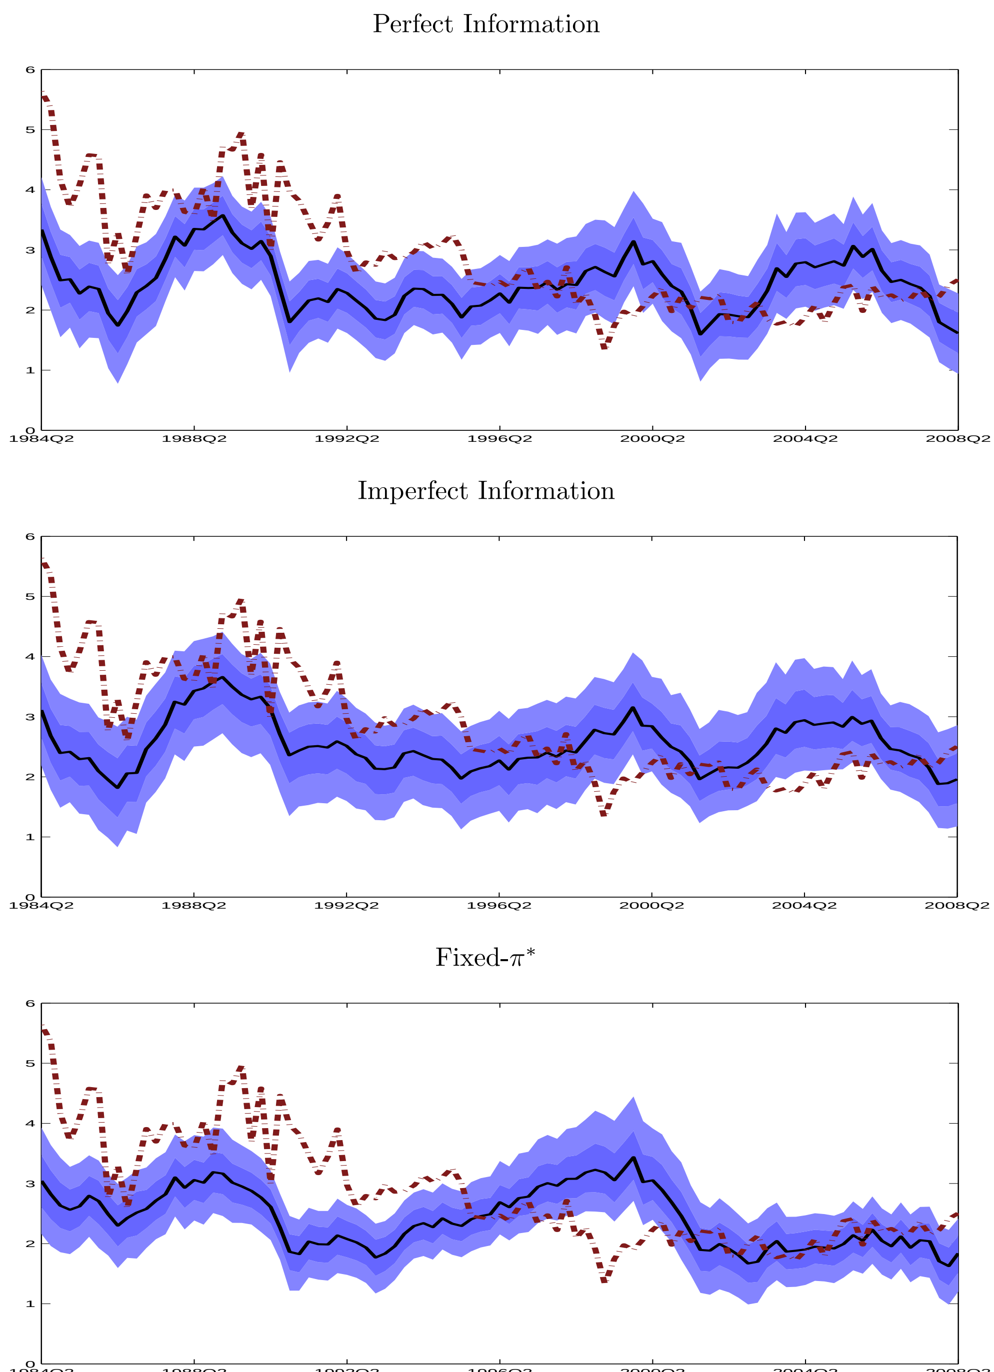

Notes: The figure shows the 4-quarters ahead median forecast for the GDP deflator (dashed and dotted line), together with 4-quarters ahead expected inflation generated by the Perfect Information (top panel), Imperfect Information (middle panel) models, and the model with a constant inflation target (Fixed-π^* – bottom panel). The predictive paths for inflation expectations are obtained as described in section 3. The solid line shows the median values and the shaded areas represent the 67 and 95% bands of the predictive distribution, respectively, computed period by period.

Figure 2: π_t^*

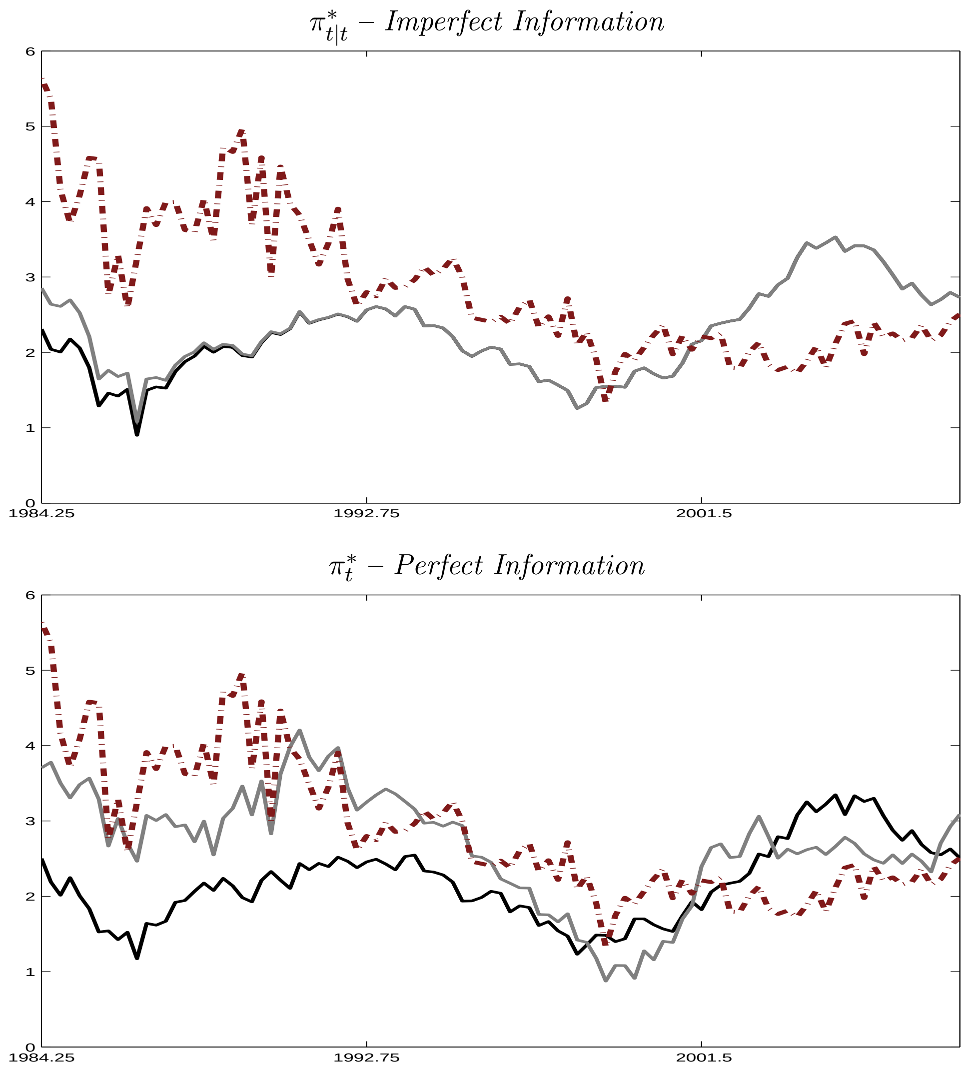

Notes: The top panel of the figure plots the mean estimate of the latent variable $\pi_{t|t}^*$ for the Imperfect Information model for the dataset without (black line) and with (gray line) inflation expectations. The bottom panel shows the mean estimate of the latent variable π_t^* for the Perfect Information model for the dataset without (black line) and with (gray line) inflation expectations. The dashed-and-dotted line shows observed inflation expectations.

Figure 3: DSGE vs VARs

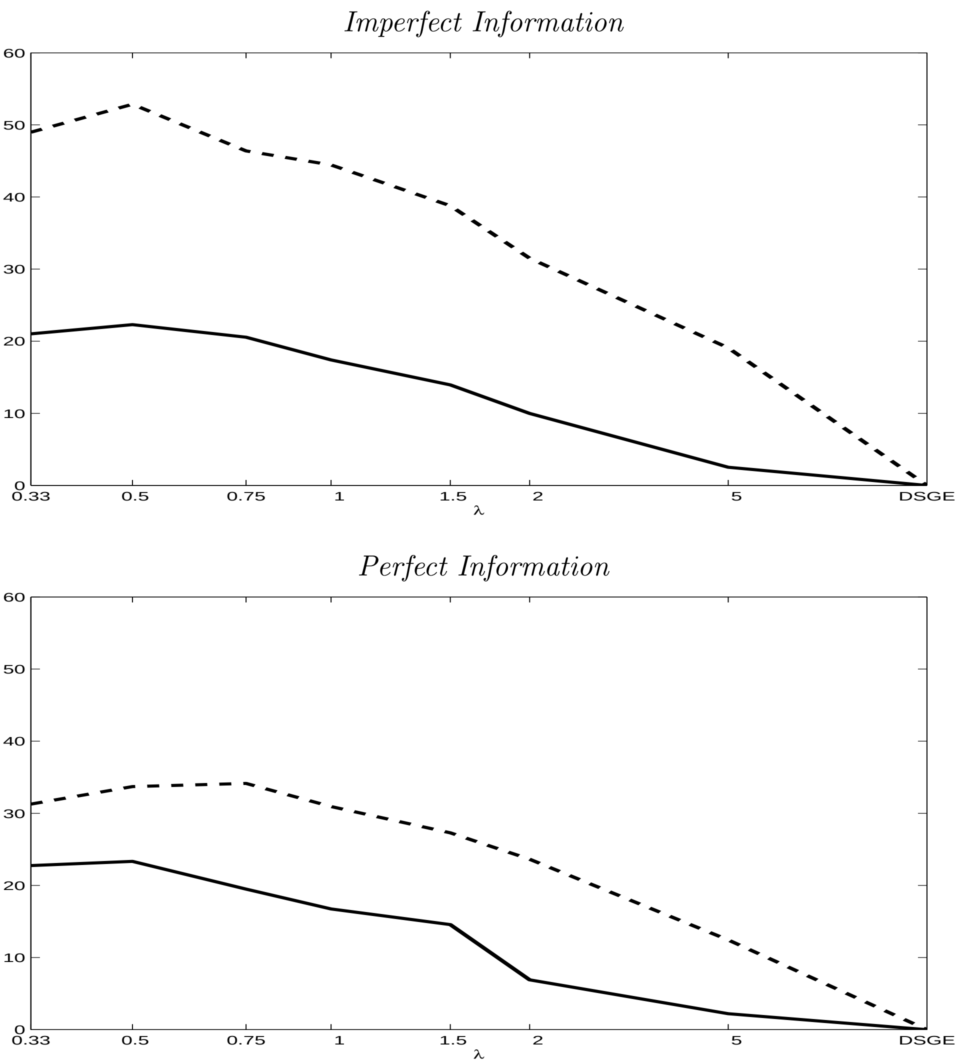

Notes: The figure plots the difference between the marginal likelihoods of the DSGE-VAR and DSGE models for the dataset with and ($ln\ p_\lambda^{VAR}(Y_{1,T}^0, Y_{1,T}^1|\mathcal{M}_i) - ln\ p^{DSGE}(Y_{1,T}^0, Y_{1,T}^1|\mathcal{M}_i)$ – dashed line) and without expectations ($ln\ p_\lambda^{VAR}(Y_{1,T}^0|\mathcal{M}_i) - ln\ p^{DSGE}(Y_{1,T}^0|\mathcal{M}_i)$ – solid line) for a grid of values for the hyper-parameter λ, which measures the tightness of the DSGE prior. The top and bottom panels present the results for the Imperfect and Perfect Information models, respectively.

Figure 4: Impulse Responses

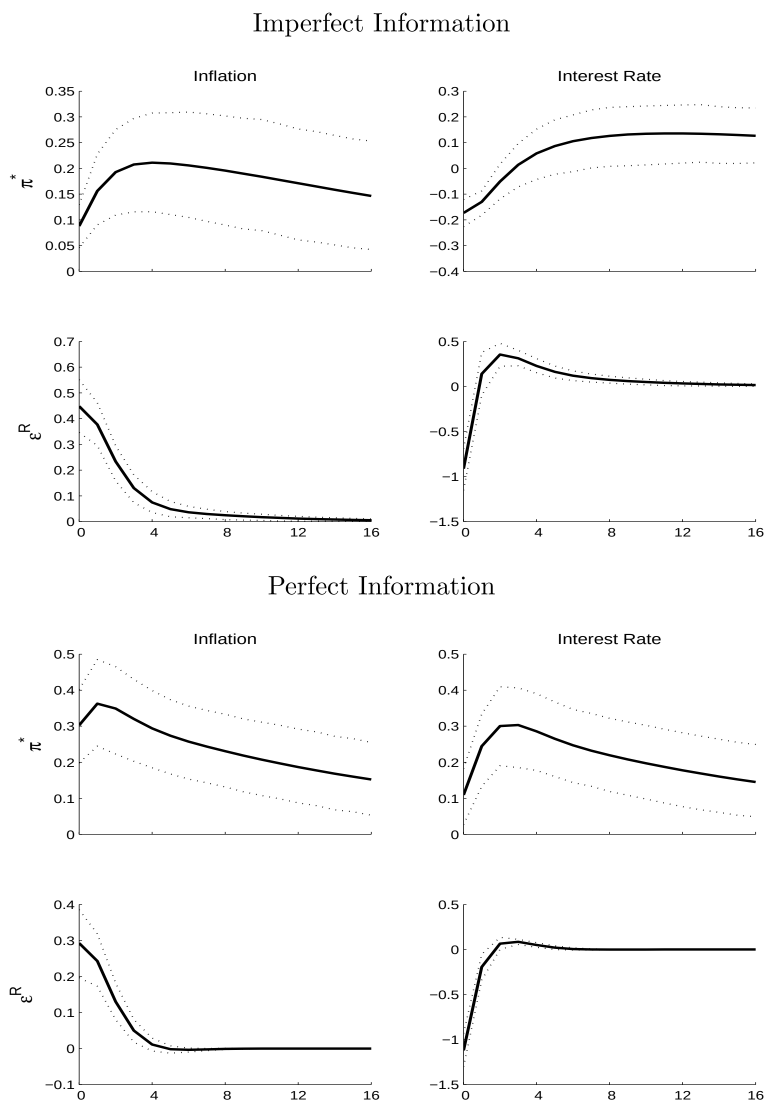

Notes: The plot show the impulse responses of inflation and the interest rate to the permanent and transitory policy shocks (ϵ_t^P and ϵ_t^R) for the Imperfect Information and the Perfect Information models obtained using the posterior from the dataset without inflation expectations.

Figure 5: Inflation Expectations: Data vs Model Prediction, 1980-2008 Sample

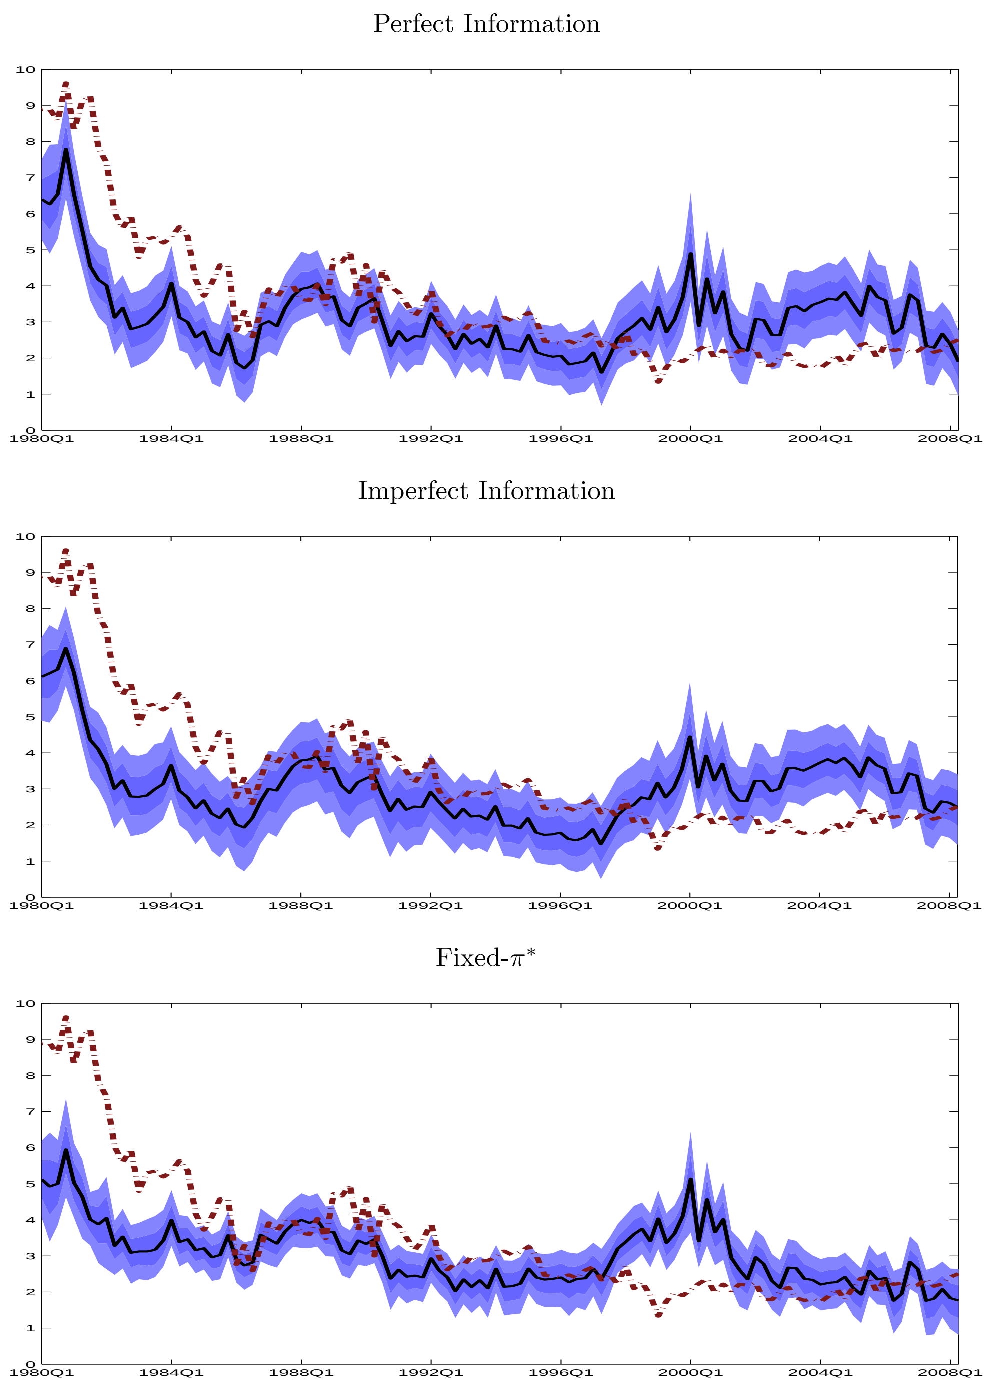

Notes: The figure shows the 4-quarters ahead median forecast for the GDP deflator (dashed and dotted line), together with 4-quarters ahead expected inflation generated by the Perfect Information (top panel), Imperfect Information (middle panel) models, and the model with a constant inflation target (Fixed-π^* – bottom panel). The predictive paths for inflation expectations are obtained as described in section 3. The solid line shows the median values and the shaded areas represent the 67 and 95% bands of the predictive distribution, respectively, computed period by period.

Table A.1: Prior Implications for Moments of the Endogenous Variables

Variables	St. Dev.			Autocorr.		
	Imperfect Information	Perfect Information	*Data*	Imperfect Information	Perfect Information	*Data*
$OutputGrowth$	3.48	3.47	*4.33*	0.39	0.39	*0.28*
$LaborSupply$	2.98	2.98	*3.20*	0.93	0.93	*0.96*
$LaborShare$	1.39	1.39	*2.24*	0.86	0.86	*0.95*
$Inflation$	3.13	3.15	*2.77*	0.71	0.72	*0.88*
$InterestRate$	4.34	4.38	*4.30*	0.85	0.85	*0.87*
$Exp.\ Inflation$	1.37	1.40		0.86	0.85	

Notes: The pre-sample statistics (column *Data*) are in italics. These statistics are computed over the sample 1959Q3-1984Q1. Inflation expectations are not available during most of the pre-sample. The in-sample standard deviation and first-order autocorrelation of inflation expectations are 1.21, and 0.86, respectively.

Figure A.1: Revisions in Inflation Data: Real Time vs Last Vintage

GDP Deflator

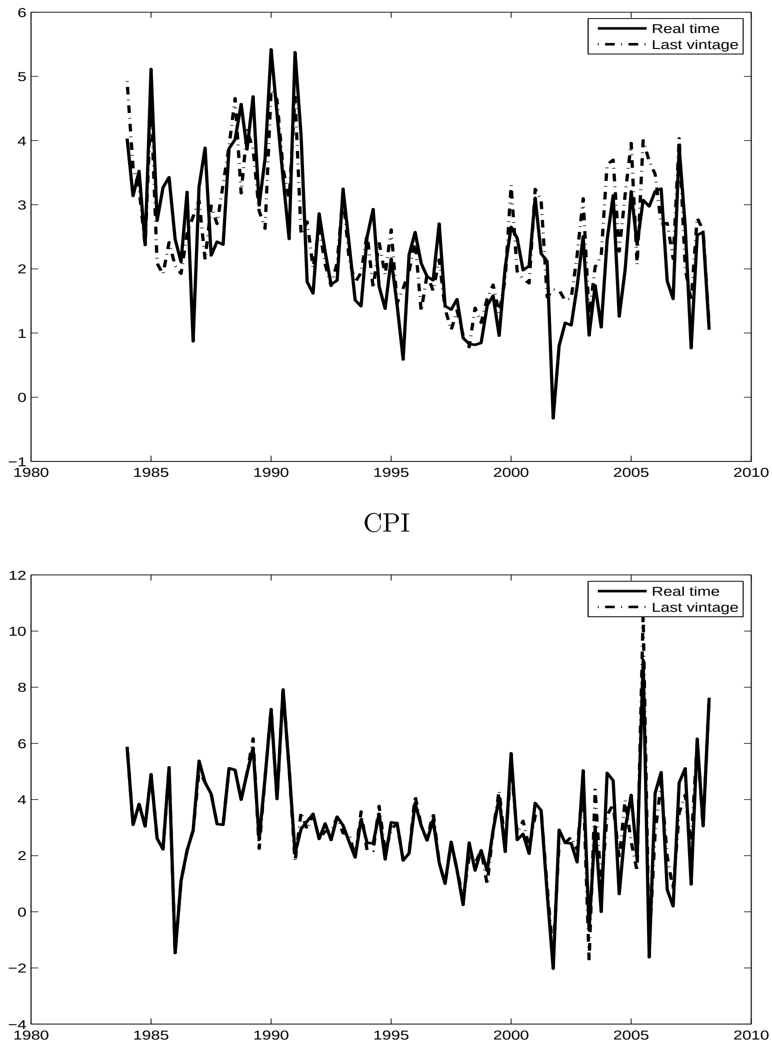

Notes: The figure plots data revisions for two measures of inflation: GDP deflator and CPI. The solid line shows the real time measure (that is, first vintage available) while the dashed-dotted line shows the most recent vintage.

Figure A.2: Impulse Responses (Dataset with Inflation Expectations)

Imperfect Information

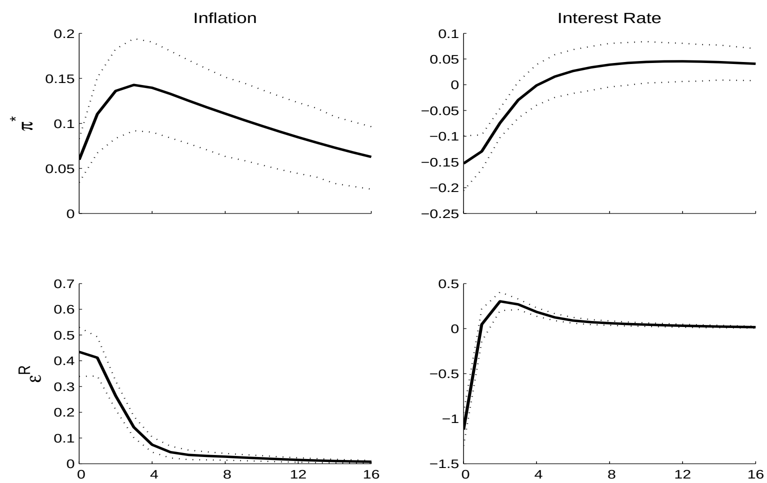

Perfect Information

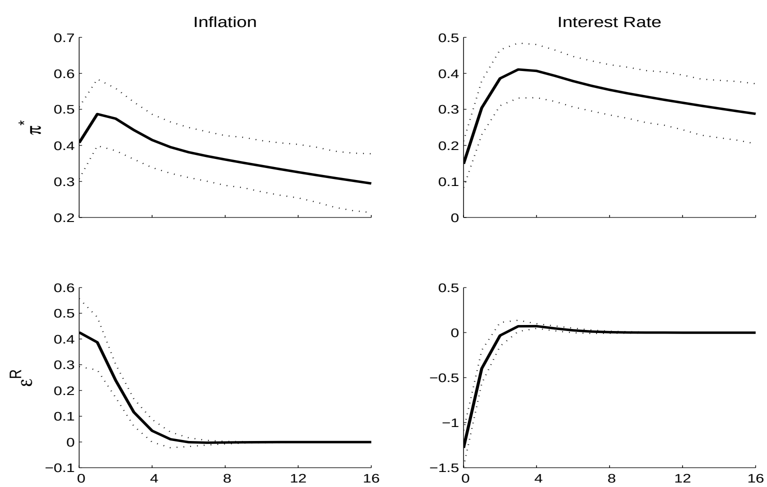

Notes: The plot show the impulse responses of inflation and the interest rate to the permanent and transitory policy shocks (ϵ_t^P and ϵ_t^R) for the Imperfect Information and the Perfect Information models obtained using the posterior from the dataset with inflation expectations.

Figure A.3: Inflation Expectations: Data vs Model Prediction, Conditioning on Initial Observation for Inflation Expectations

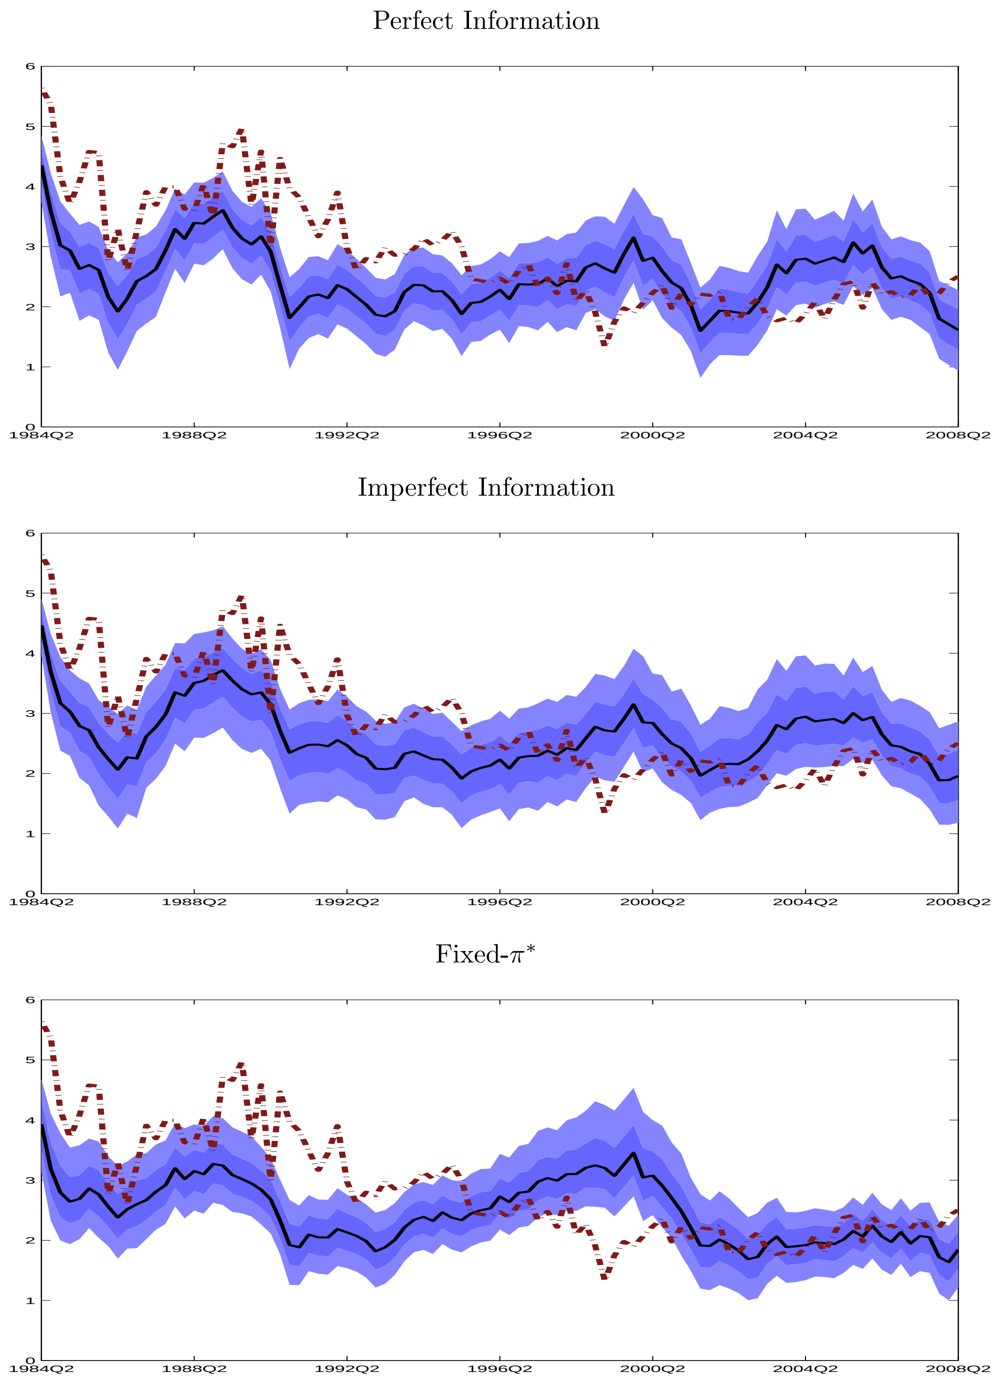

Notes: The figure shows the 4-quarters ahead median forecast for the GDP deflator (dashed and dotted line), together with 4-quarters ahead expected inflation generated by the Perfect Information (top panel), Imperfect Information (middle panel) models, and the model with a constant inflation target (Fixed-π^* – bottom panel). The predictive paths for inflation expectations are obtained as described in section 3. The solid line shows the median values and the shaded areas represent the 67 and 95% bands of the predictive distribution, respectively, computed period by period.