

EXPECTATION FORMATION AND MONETARY DSGE MODELS: BEYOND THE RATIONAL EXPECTATIONS PARADIGM

FABIO MILANI AND ASHISH RAJBHANDARI

UNIVERSITY OF CALIFORNIA, IRVINE

ABSTRACT. Empirical work in macroeconomics almost universally relies on the hypothesis of rational expectations.

This paper departs from the literature by considering a variety of alternative expectations formation models. We study the econometric properties of a popular New Keynesian monetary DSGE model under different expectational assumptions: the benchmark case of rational expectations, rational expectations extended to allow for ‘news’ about future shocks, near-rational expectations and learning, and observed subjective expectations from surveys.

The results show that the econometric evaluation of the model is extremely sensitive to how expectations are modeled. The posterior distributions for the structural parameters significantly shift when the assumption of rational expectations is modified. Estimates of the structural disturbances under different expectation processes are often dissimilar.

The modeling of expectations has important effects on the ability of the model to fit macroeconomic time series. The model achieves its worse fit under rational expectations. The introduction of news improves fit. The best-fitting specifications, however, are those that assume learning. Expectations also have large effects on forecasting. Survey expectations, news, and learning all work to improve the model’s one-step-ahead forecasting accuracy. Rational expectations, however, dominate over longer horizons, such as one-year ahead or beyond.

Keywords: Expectation Formation, Rational Expectations, News Shocks, Adaptive Learning, Survey Expectations, Econometric Evaluation of DSGE Models, Forecasting.

JEL classification: C52, D84, E32, E37, E50.

Date: This version: June 11, 2012. Revised version prepared for *Advances in Econometrics*.

We would like to thank co-editor Mark Wynne and an anonymous referee for detailed comments and suggestions, which have definitely improved the paper.

Corresponding Author: Fabio Milani, Department of Economics, 3151 Social Science Plaza, University of California, Irvine, CA 92697-5100. *Phone:* 949-824-4519. *Fax:* 949-824-2182. *E-mail:* fmilani@uci.edu. *Homepage:* <http://www.socsci.uci.edu/~fmilani>.

1. INTRODUCTION

Expectations are central to most economic decisions by households, firms, and policymakers. Also at the aggregate level, the state of expectations represents a major influence on macroeconomic outcomes.

A building block of macroeconomic theory is the assumption that all expectations in the economy are formed according to the Rational Expectations Hypothesis. Empirical work in macroeconomics is almost universally conducted under the same hypothesis. As a result, the existing evidence on the properties of microfounded models, on the values of “deep” parameters, as well as most conclusions on the reaction of the economy to different shocks, on the transmission of policy changes, and so forth, critically hinge on the validity of rational expectations as a reasonable approximation of how aggregate expectations are formed in practice.

Despite its prevalence in the profession, some economists are critical toward the rational expectations hypothesis because it endows economic agents with an extreme knowledge and capacity to process information. An alternative approach in the literature, therefore, relaxes rational expectations by assuming that agents have the same limited knowledge that a researcher estimating the model would have in real time. Agents can only observe data up to the period they live in and they use those observations to form beliefs and recursively learn about economic relationships.

Several papers have studied asymptotic convergence of systems with learning to the rational expectations equilibrium (e.g., Evans and Honkaphja, 2001, for an overview). Milani (2007, 2011) shows that learning dynamics during the transition period, instead, matters for the business-cycle behavior of macroeconomic variables.

The main scope of this paper is to provide a comprehensive assessment of how the modeling of expectations affects the properties and empirical performance of DSGE models used for monetary policy and business cycle analysis. Our setting is a baseline New Keynesian model, which has been extensively used in the monetary economics literature to study the interaction of output, inflation, and short-term interest rates.

In our analysis, we assume four different main modeling frameworks for expectations:

- (1) Rational Expectations (benchmark).
- (2) Rational Expectations, but allowing for “news” shocks.
- (3) Near-Rational Expectations with Adaptive Learning (constant gain).
- (4) Observed Survey Expectations.

Assumption (1) is the typical assumption in the literature. Agents know the structure of the model, the model parameters, and the distributions of the shocks. The source of uncertainty for agents is given by random unexpected shocks to aggregate demand, supply, and monetary policy

decisions. Assumption (2) follows a recent growing strand of literature that stresses the role of “news” as driving forces of business cycles. Economic agents still have rational expectations. But they receive news about future structural shocks. Shocks, therefore, have both an unanticipated component (as under the conventional case) and an anticipated component. Recent papers by Beaudry and Portier (2006), Jaimovich and Rebelo (2009), and Schmitt-Grohe’ and Uribe (2008) attribute to news about future technology shocks the bulk of economic fluctuations, while Milani and Treadwell (2012) single out a non-negligible role for private-sector anticipations about future monetary policy decisions. News can represent either communications by policymakers or other authorities or, simply, beliefs by the private sector about future shocks, which may or may not materialize afterwards. Assumption (3) relaxes the rational expectations assumption. Yet, the deviation is intended to be small: expectations are formed in a ‘near’-rational fashion. Economic agents use a model to form expectations that maintains the same structural form as the model solution under rational expectations. But agents lack knowledge about the model coefficients (for example, they do not know technology parameters as the Calvo price stickiness coefficient or parameters related to other consumers’ preferences). They obtain estimates of the reduced-form coefficients by using available historical data and learning about them over time. The assumption that expectations are formed near-rationally as the outcome of a learning process has been studied extensively in recent decades in a variety of settings (e.g., Evans and Honkapohja, 2001, 2013). Finally, assumption (4) chooses to exploit available data on observed, subjective, expectations from surveys (we use here expectations data from the Survey of Professional Forecasters). Those expectations are assumed to be formed from a learning model similar to the one used in (3). Economic agents, however, can deviate from the point forecasts implied by their learning model: they may be overly optimistic - by forecasting, for example, a higher future output than implied by their learning model - or overly pessimistic. These waves of over-optimism and over-pessimism are defined as ‘expectation shocks’ in the model (a similar interpretation is offered in Milani, 2011).

We study the implications of the different modeling assumptions regarding the formation of expectations on the empirical performance of the DSGE model. In particular, we investigate the impact the various expectation formation mechanisms have on:

- the posterior estimates for structural parameters;
- the estimation and characteristics of exogenous structural shocks;
- the in-sample fit of the model;
- the out-of-sample forecasting performance.

The main interest of this paper does not really lie in running a horse-race among the competing models of expectations formation and picking a winner: the paper, in fact, shows that the specifications that produce the best in-sample fit are not necessarily those that outperform the others in forecasting, and the models that do well in forecasting the short-run are different from those outperforming in the long-run. The main objective of this paper, instead, lies mostly in gauging to what extent empirical conclusions and properties of even a simple benchmark economic model, such as the New Keynesian model, may be sensitive to assumptions about the formation of expectations. We believe that showing the uncertainty surrounding the formation of expectations and its effects on the empirical results is of interest on its own. We can speculate that the modeling of expectations will have even larger effects when one moves to larger set-ups, such as medium-scale models à la Smets and Wouters (2007). It is common in the literature to almost exclusively consider rational expectations as the assumption of choice, without testing whether conclusions are robust to deviations from such hypothesis. Our scope is to show that uncertainty concerning the expectation formation mechanism should be taken seriously, since its impact on several conclusions may be substantial.

The empirical results reveal, in fact, that the modeling of expectations has large implications on the econometric properties of the model. The posterior estimates for the structural preference, technology, and policy parameters are quite sensitive to the specific way expectations are modeled. The properties of exogenous shocks also vary across specifications. For example, given the difficulty of models rational expectations to endogenously generate persistence, the estimation points toward exogenous shocks that need to evolve as AR processes with autoregressive coefficients close to one. Alternative expectation formation mechanisms may solve in part the persistence problem. The serial correlation that is required to match the persistence of observed macroeconomic variables is smaller in the model with news, and even smaller when rational expectations are replaced by either learning or survey expectations. The estimated structural shocks across model specifications display only limited correlation among each others.

The introduction of learning and survey expectations can lead to improvements in the model's in-sample ability to fit macroeconomic data. The conventional specification with rational expectations ranks last in terms of model fit, judging by the models' marginal likelihoods.

The evidence regarding out-of-sample forecasting performance is more mixed. The results differ depending on the forecast horizon. For one-quarter-ahead forecasts, survey expectations dominate the alternatives, whereas rational expectations perform less well. For longer-term forecasts, however, the rankings largely reverse: the model with rational expectations easily outperforms the

alternatives for horizons one and two-year ahead, while deviations from rational expectations given by learning and survey measures significantly worsen the forecasting performance.

2. A BENCHMARK MONETARY DSGE MODEL

This section outlines the derivation of the model equations for a typical small-scale monetary DSGE model. The derivation is now standard in the New Keynesian literature (e.g., Woodford, 2003).

We present the version of the model with the more conventional and accepted microfoundations, i.e., one which is not modified to include the so-called “mechanical” sources of persistence (e.g., Milani, 2006) as indexation to past inflation by monopolistically competitive price setters and habit formation in consumers’ preferences. The assumption of indexation has been repeatedly shown to be inconsistent with the microeconomic evidence on price setting (e.g., Nakamura and Steinsson, 2008). The evidence regarding habit formation is less clear-cut, but it seems hard to find supportive evidence using households’ consumption data (e.g., Dynan, 2000).

2.1. Households. Each household maximizes the following discounted sum of future expected utility functions

$$\max_{C,L,B} E_0 \sum_{t=0}^{\infty} \beta^t \left[e^{\tilde{g}_t} \frac{C_t^{1-\sigma^{-1}}}{1-\sigma^{-1}} - \frac{L_t^{1+\chi}}{1+\chi} \right] \quad (2.1)$$

subject to the period budget constraint

$$C_t + \frac{B_t}{P_t} = W_t L_t + \frac{(1+R_{t-1})B_{t-1}}{P_t} + \frac{D_t}{P_t} - T_t. \quad (2.2)$$

Each household derives utility from consumption C_t and disutility from hours of labor supplied L_t . The coefficient β denotes the discount factor, while σ and χ denote the elasticities of intertemporal substitution and the inverse of the elasticity of labor supply. The term $e^{\tilde{g}_t}$ denotes an aggregate taste shock. In the budget constraint (2.2), B_t denotes nominal bond holdings, P_t denotes the aggregate price level, W_t the real wage, R_t the nominal interest rate, D_t dividend distributions, and T_t net transfers or taxes.

The first order conditions imply

$$e^{\tilde{g}_t} C_t^{-\frac{1}{\sigma}} = \lambda_t \quad (2.3)$$

$$\lambda_t = \beta(1+R_t) E_t [\lambda_{t+1} (P_t/P_{t+1})] \quad (2.4)$$

$$L_t^\chi = \lambda_t W_t. \quad (2.5)$$

From (2.3) and (2.4), we can obtain the following Euler equation (after loglinearization around a zero-inflation steady state)

$$c_t = E_t c_{t+1} - \sigma (i_t - E_t \pi_{t+1} - \tilde{\rho} - E_t \Delta \tilde{g}_{t+1}), \quad (2.6)$$

where $\tilde{\rho} = -\log \beta$ is the discount rate, c_t now denotes consumption in log deviations from the steady state, i_t denotes the short-term nominal interest rate, π_t denotes the inflation rate, and $\tilde{g}_t = \log(e^{\tilde{g}_t})$.

It is possible to rewrite the Euler equation in terms of the output gap, which is the relevant variable for monetary policy, by using the resource constraint $c_t = y_t$ and the output gap definition $x_t = y_t - y_t^*$. We obtain:

$$x_t = E_t x_{t+1} - \sigma (i_t - E_t \pi_{t+1}) + g_t, \quad (2.7)$$

where the redefined demand shock g_t includes previous preference shocks \tilde{g}_t and potential output terms y_t^* .

2.2. Firms. Firms are assumed to operate under monopolistic competition. Prices are sticky à la Calvo: each firm re-optimizes its price in every period with probability $(1 - \alpha)$ and keep its price fixed to the previously set price with probability α .

Each firm maximizes the expected discounted sum of future profits to choose an optimal price p_t^*

$$\max_{p_t^*} E_t \left\{ \sum_{\tau=0}^{\infty} (\alpha\beta)^\tau \frac{\lambda_{t+\tau}}{\lambda_t} \left[p_t^* \left(\frac{p_t^*}{P_{t+\tau}} \right)^{-\theta} Y_{t+\tau} - W_{t+\tau} \left(\left(\frac{p_t^*}{P_{t+\tau}} \right)^{-\theta} \frac{Y_{t+\tau}}{A_{t+\tau}} \right)^{\frac{1}{\eta}} \right] \right\}, \quad (2.8)$$

where we have used the expressions for the product's demand curve $y_t^i = (p_t^i/P_{t+\tau})^{-\theta} Y_{t+\tau}$ and the production function $y_t^i = A_t (L_t^i)^\eta$, and where P_t denotes the aggregate price level, A_t denotes aggregate technology, θ indicates the elasticity of substitution among differentiated products, and η accounts for potentially diminishing returns to scale in the production function. Log-linearization of the problem's first order condition, along with several manipulations, leads to the familiar New Keynesian Phillips curve:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + \mu_t \quad (2.9)$$

where $\kappa \equiv \frac{(1-\alpha)(1-\alpha\beta)(\omega+\sigma^{-1})}{\alpha}$, and μ_t denotes a cost-push supply shock. The cost-push shock is often simply appended to the model in the literature (as done here), but it can also be derived endogenously by assuming a time-varying elasticity of substitution among differentiated goods θ_t , instead.

2.3. Government and Monetary Policy. Government is assumed to have access to lump-sum taxation and Ricardian equivalence holds. In this environment, the details of fiscal policy do not influence the rest of the model.

The central bank is assumed independent from government and it sets the value of a short-term interest rate, which represents its policy instrument. Monetary policy decisions are assumed to be well approximated empirically by a Taylor rule, which is described in the next section.

2.4. Monetary DSGE Model. The aggregate economy is, therefore, summarized by the following prototypical New Keynesian model, which characterizes the joint dynamics of the output gap, inflation, and the interest rate:

$$x_t = \mathbb{E}_t x_{t+1} - \sigma (i_t - \mathbb{E}_t \pi_{t+1}) + g_t \quad (2.10)$$

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + \mu_t \quad (2.11)$$

$$i_t = \rho i_{t-1} + (1 - \rho) [\chi_\pi \pi_t + \chi_x x_t + \chi_{\Delta\pi} \Delta\pi_t + \chi_{\Delta x} \Delta x_t] + \nu_t. \quad (2.12)$$

Equation (2.10) is the loglinearized consumer's Euler equation. The output gap x_t is a function of the expected one-period-ahead output gap $\mathbb{E}_t x_{t+1}$, of the ex-ante real interest rate $(i_t - \mathbb{E}_t \pi_{t+1})$, and of the demand disturbance g_t , which accounts for exogenous shift in consumers' preferences. The coefficient σ denotes the elasticity of intertemporal substitution.

Equation (2.11) is the New Keynesian Phillips curve. Current inflation is determined by expectations about future inflation $\mathbb{E}_t \pi_{t+1}$, by the current output gap x_t , and by the supply cost-push shock μ_t . The coefficient β denotes the discount factor, while κ is a composite coefficient, which moves closer to zero the higher the Calvo coefficient, and, therefore, the higher the degree of price stickiness in the economy.

Equation (2.12) is a Taylor rule, which allows for inertia in the policy instrument. The monetary policymaker sets the policy instrument i_t in reaction to movements in current inflation and current output gap with coefficients χ_π and χ_x ; as in Smets and Wouters (2003, 2007), changes in inflation and the output gap from $t - 1$ to t also enter the central bank's reaction function. The policy rate is adjusted only partially in every period toward its desired level; the coefficient ρ accounts for the degree of partial adjustment. Deviations from systematic monetary policy are captured by the term ν_t , which denotes the monetary policy shock.

All disturbances g_t , μ_t , and ν_t are assumed to follow AR(1) processes with AR coefficients ρ_g , ρ_μ , and ρ_ν .

To make clear that we focus on alternative mechanisms of expectation formation, we replace E_t in the loglinearized equations with the generic operator \mathbb{E}_t , which hence defines the expectation terms in the model. Those can either correspond to the mathematical expectation operator E_t , which stands for rational expectations, or they can instead denote subjective, non-necessarily-rational, expectations. The next section will present various alternative ways of modeling expectations.

3. THE FORMATION OF EXPECTATIONS

Expectations play a key role in the model. Households' consumption-saving and firms' price-setting decisions depend on expectations about future macroeconomic variables. While macroeconomic research typically assumes that such expectations are formed according to the rational

expectations hypothesis, here we evaluate the model under a variety of alternative expectation formation processes.

3.1. The Benchmark: Rational Expectations. The vast majority of DSGE models that are used to analyze the interaction between the dynamics of macroeconomic variables and monetary policy decisions use the rational expectations hypothesis. Economic agents are assumed to form expectations that correspond to the mathematical expectation conditional on the correct model of the economy and on knowing the values of all model parameters, the distribution of the shocks, and so forth. The source of randomness for agents remains the realization of future exogenous disturbances.

Under rational expectations, the model can be represented in state-space form as

$$\Gamma_0 \xi_t = \Gamma_1 \xi_{t-1} + \Psi w_t + \Pi \zeta_t, \tag{3.1}$$

where $\xi_t = [x_t, \pi_t, i_t, E_t x_{t+1}, E_t \pi_{t+1}, g_t, \mu_t, \nu_t]'$, $w_t = [\varepsilon_t^g, \varepsilon_t^\mu, \varepsilon_t^\nu]'$, $\zeta_t = [\zeta_t^x, \zeta_t^\pi]'$, with $\zeta_t^x = x_t - E_{t-1} x_t$ and $\zeta_t^\pi = \pi_t - E_{t-1} \pi_t$ denoting the expectational errors.

The expectational errors are uniquely determined as a function of the structural innovations (assuming that the equilibrium exists and is unique). A major advantage of rational expectations consists of removing any free parameter to characterize the formation of expectations. The cost of this approach, however, is that if expectations are not exactly formed in this way, the rational expectations hypothesis introduces a potentially sizeable misspecification, which would affect conclusions about the dynamics of macroeconomic variables, as output and inflation, the effects of policy, and the drivers of business cycle fluctuations.

The model (3.1) can be solved as in Sims (2000) to find the solution

$$\xi_t = F \xi_{t-1} + G w_t, \tag{3.2}$$

which can be joined to the measurement equation $Obs_t = H \xi_t$, linking observables to the corresponding variables in the model (through the selection matrix H), to form a linear Gaussian state-space system, which can be estimated using classical or Bayesian likelihood-based methods.

3.2. Tweaking RE: Rational Expectations with News. A recent literature has emphasized the role of news, particularly regarding future technology shocks, as potential driving forces that can give rise to expectation-driven business cycles (e.g., Beaudry and Portier, 2006).

In this environment, we maintain the assumption of rational expectations, but we extend it to include both anticipated (“news”) and unanticipated components in the shocks. The disturbances

g_t , μ_t , and ν_t now evolve as:

$$g_t = \rho_g g_{t-1} + \varepsilon_t^g + \eta_{t-h}^{g,h} \quad (3.3)$$

$$\mu_t = \rho_\mu \mu_{t-1} + \varepsilon_t^\mu + \eta_{t-h}^{\mu,h} \quad (3.4)$$

$$\nu_t = \rho_\nu \nu_{t-1} + \varepsilon_t^\nu + \eta_{t-h}^{\nu,h}. \quad (3.5)$$

where $\eta_{t-h}^{j,h}$ denotes news that becomes known in $t-h$ about shocks that will materialize h periods ahead.

A news component in the disturbances can be added in the state-space representation as follows¹ (we assume here news at horizon $h = 4$, and show only the modification due to the demand shock g_t , to save space):

$$\begin{bmatrix} g_t \\ \eta_t^{g,4} \\ \eta_{t-1}^{g,4} \\ \eta_{t-2}^{g,4} \\ \eta_{t-3}^{g,4} \end{bmatrix} = \begin{bmatrix} \rho_g & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} g_{t-1} \\ \eta_{t-1}^{g,4} \\ \eta_{t-2}^{g,4} \\ \eta_{t-3}^{g,4} \\ \eta_{t-4}^{g,4} \end{bmatrix} + \begin{bmatrix} \sigma_t^g & 0 \\ 0 & \sigma_t^{\eta_{g,4}} \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \varepsilon_t^g \\ \tilde{\eta}_t^{g,4} \end{bmatrix}. \quad (3.6)$$

where $\tilde{\eta}_t^{g,4}$ simply redefines the news shock.

In the estimation, news shocks are identified through their impact on future expectations regarding structural disturbances, which will, in turn, affect agents' economic decisions. For example, assuming again, for simplicity, a news horizon $h = 4$, we would have

$$g_t = \rho_g g_{t-1} + \varepsilon_t^g + \eta_{t-4}^{g,4} \quad (3.7)$$

$$g_{t+1} = \rho_g \left(\rho_g g_{t-1} + \varepsilon_t^g + \eta_{t-4}^{g,4} \right) + \varepsilon_{t+1}^g + \eta_{t-3}^{g,4} \quad (3.8)$$

$$E_t g_{t+1} = \rho_g g_t + \eta_{t-3}^{g,4} = \rho_g^2 g_{t-1} + \rho_g \eta_{t-4}^{g,4} + \eta_{t-3}^{g,4} \quad (3.9)$$

$$E_t g_{t+2} = \rho_g^2 g_t + \eta_{t-2}^{g,4} = \rho_g^3 g_{t-1} + \rho_g^2 \eta_{t-4}^{g,4} + \rho_g \eta_{t-3}^{g,4} + \eta_{t-2}^{g,4} \quad (3.10)$$

$$E_t g_{t+3} = \rho_g^3 g_t + \eta_{t-1}^{g,4} = \rho_g^4 g_{t-1} + \rho_g^3 \eta_{t-4}^{g,4} + \rho_g^2 \eta_{t-3}^{g,4} + \rho_g \eta_{t-2}^{g,4} + \eta_{t-1}^{g,4} \quad (3.11)$$

$$E_t g_{t+4} = \rho_g^4 g_t + \eta_t^{g,4} = \rho_g^5 g_{t-1} + \rho_g^4 \eta_{t-4}^{g,4} + \rho_g^3 \eta_{t-3}^{g,4} + \rho_g^2 \eta_{t-2}^{g,4} + \rho_g \eta_{t-1}^{g,4} + \eta_t^{g,4} \quad (3.12)$$

$$E_t g_{t+5} = \rho_g^5 g_t = \rho_g^6 g_{t-1} + \rho_g^5 \eta_{t-4}^{g,4} + \rho_g^4 \eta_{t-3}^{g,4} + \rho_g^3 \eta_{t-2}^{g,4} + \rho_g^2 \eta_{t-1}^{g,4} + \rho_g \eta_t^{g,4} \quad (3.13)$$

etc.,

which show how expectations at different horizons are affected by combinations of current and past news.

The estimation procedure, therefore, treats both the news and the unanticipated components in the shocks as unobserved, and, through the Kalman filter, it provides its best estimates of how they evolve over the sample.

¹News shocks are introduced in the same way in Schmitt-Grohe' and Uribe (2008).

An important decision in the estimation of model with news is the choice of the news horizon. Milani and Treadwell (2012) presents a comprehensive analysis of the model fit under a variety of horizons. Here, we estimate the model under horizons equal to one quarter, $h = 1$, and to four quarters, $h = 4$. We have also estimated the model with horizon $h = 8$, with similar results. Longer horizons quickly become computationally intractable, given that news substantially expands the state space.

The addition of news changes the information structure in the model. Agents still form rational expectations, but their information set is now expanded to include news and the stochastic processes of the shocks also change to include both expected and unexpected components. Hence, a comparison between the two models may not be entirely obvious. In the empirical section, we ask, however, whether these modifications are enough to change the estimates of structural parameters in the model and other model properties.

3.3. Beyond RE: Learning. The assumption of rational expectations endows agents with a vast amount of knowledge about their economic environment. A number of papers relax the assumption of rational expectations to assume, instead, adaptive learning by economic agents (e.g., Evans and Honkapohja, 2001). This literature still considers a relatively small deviation from rational expectations. Agents use a model that has the same structural form as the solution under rational expectations. But they are assumed to lack knowledge about some of the model coefficients (for example, it can be assumed that agents do not know the value of the Calvo parameter). Hence, they do not know the values of the reduced-form coefficients and they try to learn them over time using historical time series.

Under learning, we assume that agents estimate the following Perceived Law of Motion (PLM)

$$\begin{bmatrix} x_t \\ \pi_t \\ i_t \end{bmatrix} = a_t + b_t \begin{bmatrix} x_{t-1} \\ \pi_{t-1} \\ i_{t-1} \end{bmatrix} + \epsilon_t, \quad (3.14)$$

The PLM corresponds to a VAR(1) estimated on the model's endogenous variables. The VAR has the same form as the minimum state variable solution of the system under rational expectations, with the assumption, which we believe empirically realistic, that agents are unable to observe the exogenous disturbances. The coefficients in a_t and b_t are unknown to economic agents. The intercept terms in a_t have all values equal to zero under rational expectations (since variables enter (3.1)-(3.2) demeaned); this information, however, is not known to agents, who also need to learn about the values of the intercepts from historical data. Learning about a_t can be interpreted as learning about the steady states or about the trends of the variables. Given that the model that agents use remains very close to rational expectations, this approach is usually defined as *near-rational*.

The learning approach implies that economic agents do not have better information than econometricians. Under rational expectations, econometricians do not have knowledge about the model coefficients and about the realizations of innovations, while agents in the model have full knowledge. Under learning, instead, agents in the model are assumed to have a similar degree of knowledge as the econometrician estimating the model. Agents learn about coefficients in a_t and b_t using the available historical data up to each point t . They update their beliefs according to constant-gain learning as

$$\widehat{\phi}_t = \widehat{\phi}_{t-1} + \bar{\mathbf{g}}R_t^{-1}\mathbf{X}_t(\mathbf{Y}_t - \widehat{\phi}_{t-1}'\mathbf{X}_t)' \quad (3.15)$$

$$R_t = R_{t-1} + \bar{\mathbf{g}}(\mathbf{X}_t\mathbf{X}_t' - R_{t-1}) \quad (3.16)$$

where $\mathbf{Y}_t \equiv \{x_t, \pi_t, i_t\}'$, $\mathbf{X}_t \equiv \{1, \mathbf{Y}_{t-1}'\}$, and $\widehat{\phi}_t' = (a_t, b_t)$. Equation (3.15) describes the updating of beliefs regarding reduced-form coefficients, while equation (3.16) describes the updating of the associated precision matrix R_t . Agents' beliefs are, therefore, equal to their values in the previous period, plus an update in the direction of the most recent forecast error.

The constant-gain coefficient $\bar{\mathbf{g}}$, which influences the extent to which agents react to new information in every period, is estimated along with the other structural parameters in the system. By varying one parameter, the constant gain, it is possible to approximate very heterogeneous learning processes.

Expectations are formed from the PLM (3.14), using the most recent beliefs, as obtained from (3.15) and (3.16):

$$\widehat{E}_{t-1} \begin{pmatrix} x_{t+1} \\ \pi_{t+1} \\ i_{t+1} \end{pmatrix} = (I + b_{t-1})a_{t-1} + b_{t-1}^2 \begin{pmatrix} x_{t-1} \\ \pi_{t-1} \\ i_{t-1} \end{pmatrix}. \quad (3.17)$$

To break the simultaneity between the formation of expectations and the reaction of the economy (i.e., with expectations being a function of time- t variables, but also time- t variables being a function of expectations formed at time t), it is typical in the adaptive learning literature to assume that agents, when forming expectations in t about variables in $t + 1$, can observe the values of the endogenous variables only up to $t - 1$. Therefore, we denote expectations formed under learning as \widehat{E}_{t-1} . The timing, therefore, is as follows. Economic agents, at time t , run regressions from their PLM of variables in $t - 1$ on their lagged values in $t - 2$. They use the resulting beliefs and the variables they observe up to $t - 1$ to form expectations about variables in $t + 1$.

Expectations formed as in (3.17) can be substituted back into the original model (2.10)-(2.12): the resulting system is referred to as the Actual Law of Motion (ALM) of the economy. The ALM can be written in state-space form as

$$\xi_t = A_t + F_t\xi_{t-1} + Gw_t, \quad (3.18)$$

where the parameter vector A_t and the parameter matrix F_t are now time-varying due to agents' learning, and which can again be linked to the measurement equation $Obs_t = H\xi_t$ to evaluate the likelihood of the system.

3.4. Beyond RE: Observed Survey Expectations. A further departure from rational expectations consists on exploiting observed expectations from surveys for \mathbb{E}_t as variables to match in the estimation. The same strategy has been used in Milani (2011). The observed expectations are assumed to be, on average, the outcomes of a near-rational learning model. In every period, however, agents are allowed to form expectations that depart from the point forecasts implied by their near-rational learning models. They can exceed in optimism, for example by forecasting higher levels of the output gap or lower levels of inflation than suggested by the learning model, or in pessimism.

In this case, both the PLM and the constant-gain learning formulas remain given by

$$\begin{bmatrix} x_t \\ \pi_t \\ i_t \end{bmatrix} = a_t + b_t \begin{bmatrix} x_{t-1} \\ \pi_{t-1} \\ i_{t-1} \end{bmatrix} + \epsilon_t, \quad (3.19)$$

and

$$\hat{\phi}_t = \hat{\phi}_{t-1} + \bar{\mathbf{g}}R_t^{-1}\mathbf{X}_t(\mathbf{Y}_t - \hat{\phi}_{t-1}'\mathbf{X}_t)' \quad (3.20)$$

$$R_t = R_{t-1} + \bar{\mathbf{g}}(\mathbf{X}_t\mathbf{X}_t' - R_{t-1}) \quad (3.21)$$

as in the previous section.

Expectations, however, are now formed as:

$$\tilde{E}_{t-1} \begin{pmatrix} x_t \\ \pi_t \\ i_t \end{pmatrix} = a_{t-1} + b_{t-1} \begin{pmatrix} x_{t-1} \\ \pi_{t-1} \\ i_{t-1} \end{pmatrix} + \begin{pmatrix} e_t^{x_0} \\ e_t^{\pi_0} \\ 0 \end{pmatrix}, \quad (3.22)$$

and

$$\tilde{E}_{t-1} \begin{pmatrix} x_{t+1} \\ \pi_{t+1} \\ i_{t+1} \end{pmatrix} = a_{t-1} + b_{t-1}\tilde{E}_{t-1} \begin{pmatrix} x_t \\ \pi_t \\ i_t \end{pmatrix} + \begin{pmatrix} e_t^{x_1} \\ e_t^{\pi_1} \\ 0 \end{pmatrix}. \quad (3.23)$$

The variables e_t^z define the expectation shocks in the model. They are identified as the component of observed expectations that cannot be justified by the learning model. The expectation shocks can refer either to optimism and pessimism regarding future real activity ($e_t^{x_0}$ and $e_t^{x_1}$) or future inflationary pressures ($e_t^{\pi_1}$ and $e_t^{\pi_0}$), at different horizons. The expectation shocks are allowed to be serially correlated, and they evolve as AR(1) processes.

In the estimation, we will use one and two-period-ahead forecasts for output growth and inflation as additional observable variables that we need to match.

4. RATIONAL VERSUS NON-RATIONAL EXPECTATIONS ECONOMETRICS

The macroeconometrics literature has recently seen a spurt of work focused on the estimation of structural models with the use of full-information techniques, either based on maximum likelihood or, more often, on Bayesian methods.

The quasi totality of such studies impose the assumption of rational expectations. The assumption of rational expectations is typically taken as given without much analysis of how the results would differ if expectations were to only minimally deviate from fully rational.

Here we present an econometric evaluation of a popular small-scale New Keynesian model and show how many of its properties are sensitive to the way expectations are modeled.

We estimate the same model under the different expectation formation mechanisms presented in sections 3.1 to 3.4. In all cases, we use U.S. data with a sample that spans the period from 1968:IV to 2005:I (1968:IV is chosen as starting date, since the expectation data that we'll use are available from that date). We use data on output growth, calculated as the log quarterly difference of Real GDP, inflation, calculated as the quarter-to-quarter log difference in the GDP implicit price deflator, and the effective federal funds rate (transformed into a quarterly rate to be consistent with the inflation series). We assume a piecewise-linear trend in output, with a break in the slope in 1994 to capture a change in slope during the New Economy period.² We have considered estimation under a linear trend without the change in slope and the results were similar.

In the case described in section 3.4, in which expectations are proxied by survey expectations, we also use the expected one and two-period-ahead forecasts for output growth, and the expected one and two-period-ahead forecasts for inflation as additional observables that we try to match in the estimation. Given that we estimate models in which we try to infer agents' learning process, news, or excesses in optimism or pessimism over the sample, it is crucial to match as closely as possible the information set that economic subjects had available over the period. To this scope, in all cases, we use real time data in the estimation. Real time data on GDP growth and inflation are obtained from the Real Time Data Set made available by the Federal Reserve Bank of Philadelphia (for all series we use the vintage available at the time to forecasters as a description of the economy), while the effective federal funds rate, which is not revised, is obtained from FRED, the Federal Reserve

²Therefore, in the estimation, we consider an empirical proxy for the output gap, rather than the welfare-relevant definition, i.e., the deviation of output from the level that would prevail in the same economy, but under flexible prices. We use a statistical proxy here (based on a linear trend), which we keep consistent across model specifications. The model-consistent theoretical output gap would, instead, differ across specifications and, it is likely, that different assumptions about expectations would also imply diverse output gap estimates. We don't pursue that avenue here. Regarding the chosen piecewise-linear trend, the literature often assumes a break in the trend slope in the 1970s. In earlier estimations, we have used a trend with breaks in the slope both in 1973 and in 1994, but dropped the break in the 1970s in subsequent estimations, given that our sample starts only few periods in advance and also because the change in the slope seemed less apparent and less sizeable in the 1970s than in 1994 in the real-time data set that we use (the estimated trend on revised data would likely differ).

Economic Database, hosted by the Federal Reserve Bank of St. Louis. For the real-time data, we use only the first vintage of each observation, rather than the last revised vintage as common in the literature. By using only the first vintage, we therefore refrain from modeling the process of revisions. An alternative, which would be possible to consider in our state-space framework, but which would significantly complicate the analysis, would be to assume that agents observe initial releases as well as different vintages (or final vintages) of the series and impose a model for the revisions.

Observed data on output growth and inflation expectations are obtained from the Survey of Professional Forecasters. These consist of forecasts for the variables one-quarter and two-quarters ahead, and we use the mean across forecasters (we abstract from issues related to the entry and exit of forecasters in the survey, which may potentially lead to some composition bias in the sample). More details on the data and transformations that were imposed on the variables are provided in the Data Appendix section.

For each specification, we estimate the following vector of common coefficients:

$$\Theta = \{\sigma, \alpha, \rho, \chi_\pi, \chi_x, \chi_{\Delta\pi}, \chi_{\Delta x}, \rho_g, \rho_\mu, \rho_\nu, \sigma_g, \sigma_\mu, \sigma_\nu\}. \quad (4.1)$$

In the model with news, we also estimate the standard deviations of the news shocks: σ_{η_g} , σ_{η_μ} , σ_{η_ν} ; in the model with learning, we also estimate the constant gain coefficient \bar{g} and, in one case, the vector of initial beliefs $\hat{\phi}_{0|0}$; in the model with observed expectations, learning, and expectation shocks, we also estimate the constant gain coefficient, as well as the autoregressive and standard deviation coefficients ($\rho_{e_{x_1}}, \rho_{e_{x_0}}, \rho_{e_{\pi_1}}, \rho_{e_{\pi_0}}, \sigma_{e_{x_1}}, \sigma_{e_{x_0}}, \sigma_{e_{\pi_1}}$, and $\sigma_{e_{\pi_0}}$), which describe the dynamics of the expectation shocks.

We assume prior distributions that closely follow those used in recent studies that estimate comparable New Keynesian models. The prior distribution for σ , the intertemporal elasticity of substitution is a Gamma with mean 1 and standard deviation 0.5. The mean matches the values typically used in calibrated versions of general equilibrium models. The standard deviation is large enough to capture the uncertainty regarding the value of σ in the literature, given that estimates range from close to 0, particularly from microeconomic studies, to values substantially above 1, in structural macro estimations. We choose Beta distributions for the parameters that should have support between 0 and 1 from theory. The Calvo parameter has prior mean equal to 0.6 with standard deviation 0.05. This roughly matches the micro evidence on price setting, which suggests prices that remain on average fixed for 8 to 11 quarters (e.g., Nakamura and Steinsson, 2008): the prior mean is set at the low end of the estimates (but higher than the prior mean in Smets and Wouters, 2007, for example). The Beta priors for the autoregressive coefficients in the disturbance equations all have mean equal to 0.5 and standard deviation 0.15. These prior distributions remain

rather uninformative. We assume Normal distributions for the reaction coefficients in the monetary policy rule, with prior mean equal to 1.5 for the reaction to inflation and to 0.125 for the reaction to the output gap. While the priors for the Taylor rule coefficients assign a non-zero probability to regions of the parameter space that do not satisfy the Taylor principle, they are the most typical choices in the DSGE literature (indeterminacy is, instead, typically ruled out by rejecting in the MCMC procedure each draw that does not satisfy the determinacy conditions, which corresponds to assigning a zero prior to indeterminacy). An alternative would be to use Gamma priors for the Taylor coefficients. We assume inverse Gamma priors for the standard deviations of the shocks. In the models with learning, we select a prior distribution with mean 0.025 and standard deviation 0.005 for the constant gain coefficient.

The models are estimated using Bayesian methods. Draws from the posterior distribution are generated using the Metropolis-Hastings algorithm. We use the Kalman filter to evaluate the likelihood of the system at each MCMC iteration. We run 200,000 draws, discarding the first quarter of draws as initial burn-in.

5. ECONOMETRIC EVALUATION

5.1. Expectation Formation and Parameter Estimates. The posterior estimates for the structural parameters under the different expectation formation alternatives are shown in Table 1.

Assumptions about the modeling of expectations largely influence the parameter estimates. It is apparent, for example, that there is large variation in the estimate of the intertemporal elasticity of substitution across specifications, with posterior means ranging from 0.42 to 0.92. Figures 1 to 6 show the posterior distributions for a selection of the main parameters across models with different expectation formation assumptions.

First, Figures 1 and 2 clearly reveal the well known difficulty of the model under rational expectations to endogenously generate levels of persistence that can match those in macroeconomic variables. The model requires extreme degrees of serial correlation in the exogenous shocks to match such persistence. The posterior distributions for the autoregressive coefficients for the demand and supply shocks g_t and μ_t fall very close to the upper bound of one. Extending the model to allow for news about future shocks improves the model's ability to capture persistence: the posterior distributions shift to the left, with most probability mass for AR coefficients between 0.7 and 0.8. The use of either learning or survey expectations further improve the model in this direction: the updating of beliefs endogenously creates persistence, so that lower levels of serial correlation in the exogenous shocks are necessary (the posterior distributions concentrate around 0.4-0.5).

Turning to the preference, technology, and policy coefficients, we see that the values of the coefficients that are most consistent with the data seem to be sensitive to the modeling of how expectations are formed. Figure 3 shows the posterior distributions for σ , the elasticity of intertemporal substitution. The best-fitting values are relatively low under rational expectations, with mode around 0.3 and mean around 0.4. The incorporation of news, learning, or survey expectations affects the distributions, which shift to the right to assign larger probabilities to values closer to 1. Figure 6 refers to the Calvo price stickiness parameter. The posterior distribution for α falls around 0.7-0.75 under rational expectations, while the distributions move toward values around 0.85 for all other models. Figures 4 and 5 show the distributions related to the monetary policy reaction coefficients toward inflation and the output gap. The reaction toward inflation appears smaller in estimated models under learning or survey expectations, while the reaction toward the output gap is estimated to be smaller under rational expectations than under all other specifications. In most of the figures, the shifts are substantial, with posterior distributions that display small regions of overlap.

Our results here are not meant to show that one specification has to be preferred to the others. But, overall, the results provide, instead, substantial evidence that the estimates for structural parameters are far from robust to the use of alternative expectation formation mechanisms. This sensitivity to variations in the modeling of expectations is mostly ignored in the literature.

5.2. Expectation Formation and Structural Shocks. We have seen in the previous section that assumptions about expectations have important consequences on persistence and on the estimated degree of serial correlation in the structural shocks. Shocks need to be more persistent under rational expectations than they are when news, learning, or subjective expectations are allowed for.

Table 2 reports the correlation of equivalent shocks across models estimated with different expectation formation schemes. Significant differences emerge among the shocks. The demand shock estimated under rational expectations has correlation that ranges from 0.44 to 0.67 with estimated shocks in the same model, but with different expectational assumptions. The similarity between demand shocks obtained in models with news and in those with learning is also limited (the correlation is between 0.28 and 0.59). The results are along the same lines for the supply shock. There are some similarities between the shock series estimated under rational expectations and under news, but large differences with the shocks obtained in models with learning. Models with learning and survey expectations give rise to supply shocks that are similar to one another. The correlations are closer, instead, for monetary policy shocks. In this case, the shocks obtained in models with news, learning, or survey expectations are very close to each other, while there are still differences between them and the implied series under rational expectations.

As a word of caution, we would like to point out that comparing shocks across models may be difficult. For example, in the model with news, the shocks considered in the table contain both unanticipated and anticipated components, while in the other models the shock realizations are unknown to the private sector. Assumptions about the underlying stochastic processes, therefore, may somewhat differ across specifications. But it is fair to say that before assigning a clear structural interpretation to shocks, it should be shown that their dynamics are reasonably robust to different modeling assumptions, such as different sensible expectation formation mechanisms.

5.3. Expectation Formation and Model Fit. An obvious way to choose the most appropriate model of expectation formation is in terms of model fit. Table 3 reports the models' marginal likelihoods under the different expectation formation mechanisms, along with the corresponding Bayes factors, expressed by considering the model with rational expectations as the null hypothesis.³ Marginal likelihoods are a standard tool in Bayesian model comparison and have many desirable properties. For example, they automatically penalize models with additional parameters and increasing degrees of complexity (while the classical likelihood itself would only increase with complexity). Bayes factors are obtained in the table as ratios between the marginal likelihood of each model and the marginal likelihood of the benchmark model with rational expectations (assuming equal model probabilities a priori).

The model with the conventional assumption of rational expectations yields the lowest fit. Maintaining rational expectations, but extending the expectation formation process to include news about future exogenous shocks leads to improvement in fit, in particular when news refer to a longer horizon ($h = 4$ rather than $h = 1$). The models with learning, however, dominate in terms of fit. The best-fitting model assumes learning by economic agents and it allows their initial beliefs to be estimated along with the remaining parameters in the model.

The marginal likelihood for the model that uses survey data on expectations is not comparable to the others, since the set of observables is different (here, it includes also observed expectations about output growth and inflation, in addition to realized output growth, inflation, and interest rates).

Jeffrey (1961) provides an interpretative scale to judge the strength of the evidence in favor of an alternative model with respect to the model in the null hypothesis. According to his scale, the Bayes factor values in table 3 provide 'decisive' evidence for all models, except the model with news with a one-period horizon, against the rational expectations benchmark (Jeffrey indicates a value of the Bayes factor equal to 100 as cutoff, after which the evidence is considered decisive). If one had to calculate the Bayes factors of the models with learning even against the best model that

³The marginal likelihoods are calculated using Geweke's modified harmonic mean approximation.

maintains rational expectations (i.e., the model with news at horizon equal to 4 periods), the Bayes factors would still reveal ‘decisive’ evidence in favor of the learning models (Bayes factors 361.41 and 4.7572×10^4 , depending on whether initial beliefs are also estimated or not).

5.4. Expectation Formation and DSGE Forecasting Performance. In addition to in-sample fit, we evaluate the out-of-sample forecasting performance of the model regarding future output growth, inflation, and interest rates across different expectational assumptions. To generate out-of-sample forecasts, we estimate recursively each model. We start from an initial sample that spans the period between 1968:IV and 1979:IV and generate one, four, and eight-period-ahead forecasts at the end of the sample. Then, we recursively add one year of observations, by re-estimating the models for the period 1968:IV and 1980:IV, for the period 1968:IV and 1981:IV, and so forth, and each time generating the corresponding forecasts at the end of the sample.

As measure of forecasting performance we compute Root Mean Squared Errors (RMSE). We also evaluate the multivariate forecasting accuracy using the trace and the (log) determinant of the scaled forecast mean squared error matrix at each horizon h , $\Sigma_{MSE}(h)$:

$$\Sigma_{MSE}(h) = \frac{1}{N_f} \sum_{t=T}^{T+N_f-1} \bar{\epsilon}_{t+h|t} \bar{\epsilon}'_{t+h|t} \quad (5.1)$$

where N_f denotes the number of forecasts, $\bar{\epsilon}_{t+h|t} = M^{-1/2} \epsilon_{t+h|t}$, $\epsilon_{t+h|t}$ is the vector of h -step-ahead forecast errors given information in t , and M denotes a scaling matrix, which is here assumed to be diagonal with variances of the variables being forecasted on the diagonal. The determinant is usually preferred as a measure of accuracy since model rankings are invariant to the choice of the scaling matrix. Here we present both criteria, since they can, in principle, provide different information: the trace is usually affected in larger part by forecast errors for the variables that are hardest to forecast, while the determinant is influenced to a larger extent by those that are easiest to forecast (e.g., Christoffel et al., 2011).

Figure 7 displays the posterior distribution for the RMSEs relative to forecasts of output growth and inflation at one, four, and eight-period-ahead horizons, and across expectation formation models. Figure 8 shows the posterior distributions for the log determinant and trace statistics.

The model with survey expectations dominates for forecasts at the one-period horizon for output growth. The models with news and with learning both improve upon the standard case of rational expectations. Rational expectations models (in the cases with and without news) perform better for inflation forecasting: the RMSEs are smaller than those obtained under learning. Moreover, the model with rational expectations performs extremely well over longer horizons.

The measures of multivariate forecast accuracy indicate that survey expectations carry information that allows the econometrician to improve the one-period-ahead forecasting performance of

the model. The model with survey expectations ranks first, followed by the model with news and learning; the model with rational expectations ranks last. In this case, therefore, the rankings based on in-sample model fit and out-of-sample forecasting performance are similar. Moving to longer forecasting horizons, however, leads to major reversals in the rankings. At the four and eight period horizons, the model with rational expectations dominates all other alternatives. The model with rational expectations extended to include news performs slightly worse and comes second in the ranking. The model with survey expectations and, particularly, the model with learning appear to perform rather poorly if evaluated in terms of their forecasting ability over long horizons.

While a number of papers have focused on assessing the forecasting success of DSGE models under rational expectations, the evidence is scant under learning. A recent paper that, among other things, seeks to assess the forecasting ability of a model with learning is Slobodyan and Wouters (2011). That paper seems to reach a similar conclusion: learning helps in forecasting in the short run, but it is outperformed by rational expectations over longer horizons.

Identifying the reasons why rational expectations perform well in the long run requires additional work. Here, we point out that rational expectations enjoy some advantages over the alternatives that may be fruitful in forecasting. For example, the learning models assume that agents do not have information regarding the steady states of the variables and that they attempt to learn them over time in the same way as they learn about other parameters (i.e., they learn about coefficients in the vector a_t in (3.14), which can be interpreted as steady-state coefficients). The steady-states are, instead, perfectly known at all times under rational expectations. This assumption of full knowledge about the steady states may favor the model with rational expectations over the medium to long run. Moreover, the model with rational expectations assumes that agents dispose of information up to time t when forming expectations in $t + 1$ and that they observe the values and histories of the shocks in time t . Under learning, we assume that agents in t can observe only endogenous variables up to $t - 1$ and that they do not know the values of the shocks. This limited knowledge may again favor the model with rational expectations.

Another feature that may be penalizing the learning models at longer horizons is the possibility that the agents' real-time PLMs at least in some periods and in some draws may become unstable. In the estimation, we don't impose stability or other constraints on agents' beliefs (letting them be exclusively driven by the data) and it's possible that in some situations the PLMs are characterized by complex (or unstable, although this doesn't happen in the full-sample estimation) eigenvalues. Instability would clearly worsen the forecasting performance, especially at longer horizons. A more in-depth investigation is beyond the scope of this paper, but it will be interesting in future research

to check to what extent adding some degree discipline to agents' beliefs would improve the long-run forecasting performance of learning models.

Finally, and more from a modeling perspective, the model with learning is imposed on the same loglinearized model equations that are obtained under rational expectations; current macroeconomic variables depend exclusively on expectations up to $t + 1$. Preston (2005), however, shows that the derivation of the model under subjective expectations can lead to laws of motion that imply that current values of the variables depend on long-horizon expectations as well. A model with learning with long-horizon expectations has been estimated in Milani (2006). Maybe extending the learning model to allow for long-horizon expectations may improve its forecasting performance beyond the one-period-ahead horizons.

6. WHAT HAVE WE LEARNED?

The main scope of the paper was to investigate whether estimates of structural parameters and the econometric properties of macroeconomic models were likely robust to some popular deviations from the benchmark hypothesis of rational expectations.

The empirical results we found carry several messages. First, it is unrealistic to assume that conclusions that are obtained for DSGE model, even in small scale environments, are likely to be reasonably robust to even small deviations from the benchmark assumption of rational expectations.

The previous section has shown that parameter estimates, estimates of the unobserved structural shocks, in-sample model fit, and out-of-sample forecasting performance are largely affected by the modeler's assumptions about how expectations are formed.

More work is, instead, needed to choose among the various models of expectations formation.

The rational expectations hypothesis, which is used in the vast majority of theoretical and empirical work in macroeconomics, may worsen the fit of the model. The conventional model with rational expectations, in fact, ranks last in model fit according to the marginal likelihoods. The model with learning fits the data substantially better than do the alternatives. If one is interested in forecasting one-quarter-ahead macroeconomic variables, the use of survey expectations is valuable. The model with survey expectations achieves the overall best out-of-sample forecasting performance in the very short run. But when the attention turns to slightly longer horizons (from one-year ahead and beyond), the model with rational expectations strongly outperforms the alternatives.

These results may be taken to suggest that appropriate models of expectations formation may be adjusted depending on the purpose at hand. Researchers interested in modeling short-run economic dynamics may need to take economic agents' learning processes and less-than-fully rational expectations into account, while researchers interested in economic adjustment in the medium to long-run may do well or better by maintaining the assumption of rational expectations.

Our main point, however, is that, at a minimum, the paper’s results should suggest the need for modelers and econometricians to check the sensitivity of their results to alternative expectation formation processes.

7. CONCLUSIONS

Current macroeconomic theory is built on the assumption that economic agents form expectations according to the rational expectation hypothesis. The conclusions that are derived from empirical work in macroeconomics also hinge on the validity of rational expectations.

In this paper, we have evaluated the consequences of relaxing the assumption of rational expectations in a popular small-scale monetary DSGE model in the New Keynesian tradition. We have shown that the econometric properties of the model are extremely sensitive to the way expectations are modeled. The formation of expectations is a dimension in which the effects of misspecification and any analysis of sensitivity are largely lacking in the literature. The evidence presented in this paper shows that the formation of expectations should be studied more critically in estimated DSGE models.

Limitations. Besides the various expectation models that we examined in this paper, we recognize that there are other alternatives that have not been included. For example, it can be assumed that economic agents retain rational expectations, but that their information is sticky or that they optimally choose to be at times inattentive to economic developments (Mankiw and Reis, 2007, Sims, 2010). Within learning models, there are various other ways in which agents could be assumed to learn, which may differ from the constant gain learning assumed in this paper (e.g., recursive-least-squares learning, or learning with a time-varying endogenous gain, as in Milani, 2008). Finally, they may have imperfect information about some of the shocks or variables and they may solve signal extraction problems (e.g., Levine et al., 2010).

Moreover, we would like to point out several potential limitations regarding the perspective offered in this paper.

First, we have employed a rather narrow interpretation of rational expectations: here agents have full knowledge about the structure of the model, its parameters, the distributions of the shocks, and the history of endogenous variables and shocks up to each period in time. The definition of rational expectations, however, may be broadened (and, indeed, it has often been broadened in the literature) to simply mean that agents will rationally use all the information that is available to them at each point in time: rational expectations can, therefore, coexist with imperfect knowledge, such as imperfect information about some of the shocks (the literature has often considered imperfect information about a time-varying Fed’s inflation target, for example), learning about some parameters (e.g., unknown, and possibly time-varying, coefficients in the Taylor rule), and so forth.

In such cases, the information sets attributed to agents under rational expectations and adaptive learning are brought closer together and separating the two may possibly become a much harder task. The main difference would still remain that in models with learning such as those used in this paper, agents are not fully rational: given incomplete knowledge, they do not learn optimally given the assumed model of the economy, while they would optimally process the information under rational expectations. Near-rational models with learning are meant to provide an approximation of optimal solutions that may be too difficult to compute by agents in situations in which the amount of knowledge is limited.

Second, the model that we have chosen as reference may be misspecified. In the paper, we are implicitly always considering a “joint hypothesis”: one regarding the form of the expectation formation mechanism and the other implicitly assuming that the structural model we use is an appropriate description of the data-generating process for the macroeconomic series we try to explain. If the structural model is misspecified, structural estimates may be biased, we may not be able to fully recover the structural shocks, and the biases may translate into different conclusions regarding the expectation-formation mechanism. At the same time, we believe that this is true in most cases in the literature. In the empirical DSGE literature, for example, it is typical to test the role of various frictions (for example, by shutting down in turn habit formation, adjustment costs in investment, price or wage stickiness, etc.) or extensions to the baseline model (for example, by adding a financial accelerator mechanism or search and matching in the labor market): tests of their importance are also implicitly joint test of such model features, but also of the expectation formation hypothesis. This paper suggests that this practice may be problematic: assumptions about expectations may easily be as or more important than assumptions about other details in the model.

Third, the comparison across the four models with different expectation schemes may not always be straightforward, given differences in some of the auxiliary modeling assumptions. For example, under learning we have assumed that agents dispose of an information set including variables up to $t - 1$, while under rational expectations their information set includes variables up to t ; in the benchmark rational expectations model and in the models with learning, the shocks are entirely unanticipated by the private sector, while in the models with news, the shocks contain both unanticipated and anticipated components. It will be worthwhile testing the extent to which these differences matter.

Finally, a pervasive problem in the the DSGE literature lies in the choices regarding how to detrend real variables, in this case GDP, and how to construct measures of the output gap. Most papers typically proceed by either detrending the data before the estimation using a statistical

procedure (linear or quadratic trend, Hodrick-Prescott filter, band-pass filter, and so forth), or using a model-consistent output gap, calculated as the deviation between actual output and the corresponding level of output that would obtain in the same exact economy, but under flexible prices. In this paper, we have chosen to consider the simplest case of a linear trend (although with a break in slope), to keep the trend as consistent as possible across model specifications, and with the idea that this would be more likely to approximate the detrending procedure forecasters had in mind, at least over large portions of the sample, when forming their expectations. But it is well known that different detrending schemes and the use of statistical measures in place of the theoretically-consistent output gap may have important effects on the results (e.g., Neiss and Nelson, 2003, 2005). A more detailed study of the interaction between detrending procedures, output gap measures, and expectation formation assumptions, is certainly warranted.

The aim of this paper was to shed light on some of the empirical implications of alternative expectation formation schemes in a benchmark monetary business cycle model. But future research is needed to provide more definitive evidence on the best model of expectations formation and on the role of alternative expectational assumptions in more complicated models.

A. DATA APPENDIX

We document here the observable variables that we use in the estimation, the transformations that were imposed on the series, and the corresponding variable in the model. The data are all for the U.S. and the sample is 1968:IV-2005:I.

x_t : Real GDP: we use Real GDP (acronym: ROUTPUT), Billions of Real Dollars, Seasonally Adjusted, obtained from the Real Time Data Set for Macroeconomists hosted by the Federal Reserve Bank of Philadelphia. We compute quarterly output growth as $[\log(\text{Real GDP}_t) - \log(\text{Real GDP}_{t-1})] \times 100$ and link output growth to x_t in the model through the observation equation: $[\log(\text{Real GDP}_t) - \log(\text{Real GDP}_{t-1})] \times 100 = \Delta x_t + \gamma_t$, where γ_t capture the piecewise linear trend, i.e. $\gamma_t = \gamma_1$ before 1994 and $\gamma_t = \gamma_2$ after 1994.

π_t : Inflation rate: to construct quarterly inflation, we use the Implicit Price Deflator series (acronym: P), index level, seasonally adjusted, obtained from the Real Time Data Set for Macroeconomists hosted by the Federal Reserve Bank of Philadelphia. The base year in the real-time vintages vary, therefore, we re-express all observations using the same base year. Inflation is calculated as $[\log(P_t) - \log(P_{t-1})] \times 100$. The series is demeaned. The observation equation is simply: $[\log(P_t) - \log(P_{t-1})] \times 100 = \pi_t$.

i_t : Nominal Interest rate: we use the Effective Federal Funds Rate (acronym: FEDFUNDS) from FRED, the Federal Reserve Database maintained by the Federal Reserve Bank of St. Louis. The Federal Funds rate is divided by four to express it in quarterly rates. The observable is matched to the variable i_t in the model as $FFR_t/4 = i_t$.

In the specification that uses observed expectations, the estimation adds to the previous set of observables the following:

$\widehat{E}_{t-1}x_t, \widehat{E}_{t-1}x_{t+1}$: GDP growth forecasts (one-period-ahead and two-period-ahead): we use forecasts for real GDP (acronym: RGDP) obtained from the Survey of Professional Forecasters. We use mean responses across forecasters as our expectations measures. We use columns four and five corresponding to forecasts RGDP2 and RGDP3 (one-quarter and two-quarter ahead).

Forecasts for the growth rates can be computed as $[\log(\text{RGDP2}_t) - \log(\text{RGDP1}_t)] \times 100$ and $[\log(\text{RGDP3}_t) - \log(\text{RGDP2}_t)] \times 100$, where RGDP1 represents the forecasters' estimate in t of real GDP in $t - 1$. The forecasts are matched to the expectations in the model using the observation equations $[\log(\text{RGDP2}_t) - \log(\text{RGDP1}_t)] \times 100 = \widehat{E}_{t-1}\Delta x_t + \widehat{\gamma}_t$ and $[\log(\text{RGDP3}_t) - \log(\text{RGDP2}_t)] \times 100 = \widehat{E}_{t-1}\Delta x_{t+1} + \widehat{\gamma}_t$, where we allow the trend coefficients $\widehat{\gamma}_t$ inferred by forecasters to differ from those in γ_t .

$\widehat{E}_{t-1}\pi_t, \widehat{E}_{t-1}\pi_{t+1}$: Inflation rate (one-period-ahead and two-period-ahead): we use forecasts for the GDP implicit price deflator (acronym: PGDP) obtained from the Survey of Professional Forecasters. We use mean responses across forecasters as our expectations measures. We use columns four and five corresponding to forecasts PGDP2 and PGDP3 (one-quarter and two-quarter ahead).

Forecasts for inflation rates are computed as $[\log(PGDP2_t) - \log(PGDP1_t)] \times 100$ and $[\log(PGDP3_t) - \log(PGDP2_t)] \times 100$, where PGDP1 represents the forecasters' estimate in t of the price level in $t-1$. The expected inflation series are demeaned before the estimation. The forecasts are matched to the expectations in the model using the observation equations $[\log(PGDP2_t) - \log(PGDP1_t)] \times 100 = \widehat{E}_{t-1}\pi_t$ and $[\log(PGDP3_t) - \log(PGDP2_t)] \times 100 = \widehat{E}_{t-1}\pi_{t+1}$.

REFERENCES

- [1] Paul Beaudry and Franck Portier, 2006. "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, vol. 96(4), pages 1293-1307.
- [2] Kai Christoffel, Gnter Coenen, and Anders Warne, 2010. "Forecasting with DSGE models," Working Paper Series 1185, European Central Bank.
- [3] Karen E. Dynan, 2000. "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*, vol. 90(3), pages 391-406.
- [4] George W. Evans and Seppo Honkapohja, 2001. *Learning and Expectations in Economics*. Princeton University Press, Princeton, NJ.
- [5] George W. Evans and Seppo Honkapohja, 2013. Learning as a Rational Foundation for Macroeconomics and Finance, (with Seppo Honkapohja), forthcoming in *Rethinking Expectations: The Way Forward for Macroeconomics*, Roman Frydman and Edmund S. Phelps (eds.), Princeton University Press.
- [6] Nir Jaimovich and Sergio Rebelo, 2009. "Can News about the Future Drive the Business Cycle?," *American Economic Review*, vol. 99(4), 1097-1118.
- [7] Harold Jeffrey, 1961. *Theory of Probability*, Oxford, U.K.: Oxford University Press.
- [8] Paul Levine, Joseph Pearlman, George Perendia, and Bo Yang, 2010. "Endogenous Persistence in an Estimated DSGE Model under Imperfect Information," CDMA Working Paper Series 1002.
- [9] Gregory N. Mankiw, and Ricardo Reis, 2007. "Sticky Information in General Equilibrium," *Journal of the European Economic Association*, vol. 5(2-3), pages 603-613.
- [10] Fabio Milani, 2006. "A Bayesian DSGE Model with Infinite-Horizon Learning: Do "Mechanical" Sources of Persistence Become Superfluous?," *International Journal of Central Banking*, vol. 2(3), pages 87-106.
- [11] Fabio Milani, 2007. "Expectations, Learning and Macroeconomic Persistence," *Journal of Monetary Economics*, vol. 54(7), pages 2065-2082.
- [12] Fabio Milani, 2008. "Learning and Time-Varying Macroeconomic Volatility", mimeo, University of California, Irvine.
- [13] Fabio Milani, 2011. "Expectation Shocks and Learning as Drivers of the Business Cycle", *Economic Journal*, Vol. 121, Iss. 552, pages 379-401.
- [14] Fabio Milani and John Treadwell, 2012. "The Effects of Monetary Policy "News" and "Surprises"", forthcoming, *Journal of Money, Credit and Banking*.
- [15] Emi Nakamura and Jon Steinsson, 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, vol. 123(4), pages 1415-1464.
- [16] Katharine S. Neiss, and Edward Nelson, 2003. "The Real-Interest-Rate Gap as an Inflation Indicator," *Macroeconomic Dynamics*, 7(2), 239-62.
- [17] Katharine S. Neiss, and Edward Nelson, 2005. "Inflation Dynamics, Marginal Cost, and the Output Gap: Evidence from Three Countries," *Journal of Money, Credit, and Banking*, 37(6), 1019-45.
- [18] M. Hashem Pesaran, 1987. *The Limits to Rational Expectations*, Oxford and New York: Blackwell Publishing.
- [19] Bruce Preston, 2005. "Learning about Monetary Policy Rules when Long-Horizon Expectations Matter," *International Journal of Central Banking*, vol. 1(2).
- [20] Stephanie Schmitt-Grohe' and Martin Uribe, 2008, "What's News in Business Cycles?," NBER Working Paper 14215.
- [21] Christopher A. Sims, 2000. "Solving Linear Rational Expectations Models", *Computational Economics*, 20, 1-20.
- [22] Christopher A. Sims, 2010. "Rational Inattention and Monetary Economics", *Handbook of Monetary Economics*, vol. 3, 155-181, eds. Benjamin M. Friedman and Michael Woodford, Amsterdam: Elsevier.
- [23] Sergey Slobodyan and Raf Wouters, 2009. "Learning in an Estimated Medium-Scale DSGE Model," CERGE-EI Working Papers wp396.
- [24] Frank Smets and Rafael Wouters, 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, vol. 1(5), 1123-1175.
- [25] Frank Smets and Rafael Wouters, 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, vol. 97(3), pages 586-606.
- [26] Michael Woodford, 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton, NJ: Princeton University Press.

Params.	Prior distr.			Posterior distributions					
	Distr.	Mean	S.d.	RE	News $h = 1$	News $h = 4$	Learning	Learning I.B.	Survey
σ	G	1	0.5	0.42 [0.25,0.62]	0.59 [0.33,0.98]	0.83 [0.51,1.18]	0.92 [0.50,1.33]	0.78 [0.47,1.11]	0.90 [0.54,1.29]
α	B	0.6	0.05	0.75 [0.68,0.81]	0.83 [0.79,0.88]	0.86 [0.81,0.90]	0.84 [0.81,0.88]	0.82 [0.77,0.86]	0.83 [0.79,0.87]
ρ	B	0.7	0.1	0.62 [0.50,0.73]	0.80 [0.73,0.85]	0.82 [0.77,0.86]	0.89 [0.85,0.93]	0.89 [0.84,0.93]	0.90 [0.85,0.93]
χ_π	N	1.5	0.125	1.66 [1.42,1.90]	1.78 [1.55,1.97]	1.73 [1.51,1.94]	1.42 [1.17,1.67]	1.43 [1.21,1.67]	1.44 [1.20,1.68]
χ_x	N	0.125	0.0625	0.01 [-0.03,0.05]	0.18 [0.07,0.30]	0.24 [0.16,0.33]	0.23 [0.12,0.35]	0.22 [0.11,0.33]	0.23 [0.11,0.35]
$\chi_{\Delta\pi}$	N	0.1	0.05				0.08 [-0.02,0.17]	0.08 [0,0.16]	0.09 [-0.01,0.17]
$\chi_{\Delta x}$	N	0.1	0.05				0.10 [0.02,0.20]	0.09 [0,0.19]	0.08 [-0.01,0.17]
ρ_g	B	0.5	0.15	0.83 [0.76,0.90]	0.73 [0.62,0.82]	0.72 [0.60,0.84]	0.52 [0.38,0.66]	0.55 [0.41,0.68]	0.51 [0.35,0.66]
ρ_μ	B	0.5	0.15	0.99 [0.98,0.99]	0.92 [0.79,0.99]	0.76 [0.60,0.88]	0.42 [0.23,0.63]	0.50 [0.31,0.69]	0.48 [0.32,0.64]
ρ_ν	B	0.5	0.15	0.42 [0.30,0.53]	0.26 [0.14,0.40]	0.23 [0.11,0.34]	0.27 [0.14,0.41]	0.25 [0.12,0.41]	0.28 [0.15,0.43]
σ_g	IG	0.33	1	0.11 [0.07,0.19]	0.14 [0.08,0.22]	0.28 [0.20,0.38]	0.99 [0.86,1.13]	0.97 [0.86,1.08]	0.86 [0.76,0.96]
σ_μ	IG	0.33	1	0.42 [0.24,0.66]	0.12 [0.08,0.18]	0.11 [0.08],0.15	0.40 [0.34,0.45]	0.40 [0.35,0.45]	0.33 [0.30,0.37]
σ_ν	IG	0.25	1	0.29 [0.23,0.36]	0.13 [0.08,0.21]	0.08 [0.06,0.11]	0.21 [0.18,0.23]	0.20 [0.18,0.23]	0.21 [0.18,0.23]
σ_η^g	IG	0.25	1		0.35 [0.16,0.60]	0.37 [0.17,0.55]			
σ_η^μ	IG	0.25	1		0.09 [0.06,0.16]	0.13 [0.08,0.21]			
σ_η^ν	IG	0.25	1		0.17 [0.11,0.22]	0.21 [0.18,0.24]			
\bar{g}	B	0.025	0.005				0.023 [0.015,0.033]	0.018 [0.011,0.028]	0.025 [0.017,0.033]
ρ_e^{x1}	B	0.5	0.15						0.56 [0.43,0.69]
$\rho_e^{\pi1}$	B	0.5	0.15						0.48 [0.33,0.63]
ρ_e^{x0}	B	0.5	0.15						0.41 [0.27,0.56]
$\rho_e^{\pi1}$	B	0.5	0.15						0.65 [0.56,0.74]
σ_e^{x1}	IG	0.33	1						0.95 [0.83,1.06]
$\sigma_e^{\pi1}$	IG	0.33	1						0.16 [0.14,0.18]
σ_e^{x0}	IG	0.33	1						0.20 [0.18,0.23]
$\sigma_e^{\pi0}$	IG	0.33	1						0.21 [0.17,0.24]

TABLE 1. Prior distributions and posterior estimates, across expectation formation models. *Note:* The numbers below in the table denote posterior mean; the numbers in square brackets denote 95% posterior probability intervals.

	RE	News h=1	News h=4	Learning	Learning I.B.	Survey
<u>Demand Shock g_t</u>						
RE	1					
News h=1	0.67	1				
News h=4	0.59	0.92	1			
Learning	0.52	0.28	0.33	1		
Learning I.B.	0.58	0.35	0.38	0.99	1	
Survey	0.46	0.44	0.59	0.66	0.66	1
<u>Supply Shock μ_t</u>						
RE	1					
News h=1	0.81	1				
News h=4	0.76	0.98	1			
Learning	0.28	0.53	0.57	1		
Learning I.B.	0.31	0.53	0.56	0.99	1	
Survey	0.40	0.58	0.60	0.83	0.83	1
<u>MP Shock ν_t</u>						
RE	1					
News h=1	0.88	1				
News h=4	0.80	0.99	1			
Learning	0.61	0.90	0.94	1		
Learning I.B.	0.67	0.93	0.96	0.99	1	
Survey	0.66	0.92	0.95	0.99	0.99	1

TABLE 2. Estimated structural shocks: correlation across expectation formation models.

	Marginal Likelihood	Bayes Factor
RE	-310.39	1
News $h = 1$	-309.81	1.79
News $h = 4$	-296.83	7.7452×10^5
Learning	-290.94	2.7992×10^8
Learning I.B.	-286.06	3.6846×10^{10}
Survey Exp.	-	-

TABLE 3. Bayesian Model Comparison: models' Log Marginal Likelihoods and Bayes Factors. *Note:* The marginal likelihood for the model with survey expectations is not reported as not comparable with the others (given that the estimation uses a different set of observables). According to Jeffrey's (1961) scale of evidence, Bayes factors above values of 100 represent 'decisive' evidence in favor of a model versus the other.

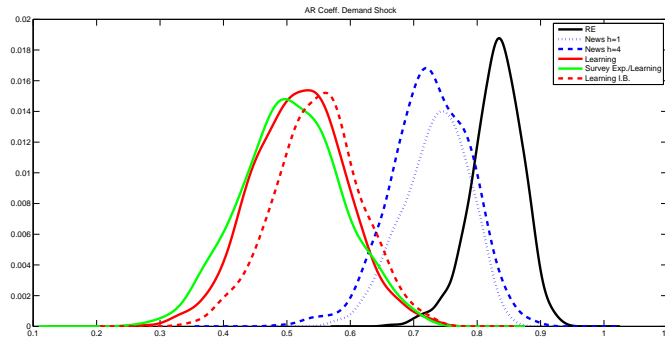


FIGURE 1. Posterior distributions: AR coefficient for demand shock g_t . *Note:* the figure shows posterior distributions in the estimated models across different expectation formation schemes.




FIGURE 2. Posterior distributions: AR coefficient for supply shock μ_t .




FIGURE 3. Posterior distributions: intertemporal elasticity of substitution σ .




FIGURE 4. Posterior distributions: monetary policy reaction toward inflation χ_π .




FIGURE 5. Posterior Distributions: monetary policy reaction toward output gap χ_x .




FIGURE 6. Posterior distributions: Calvo Price Stickiness α .

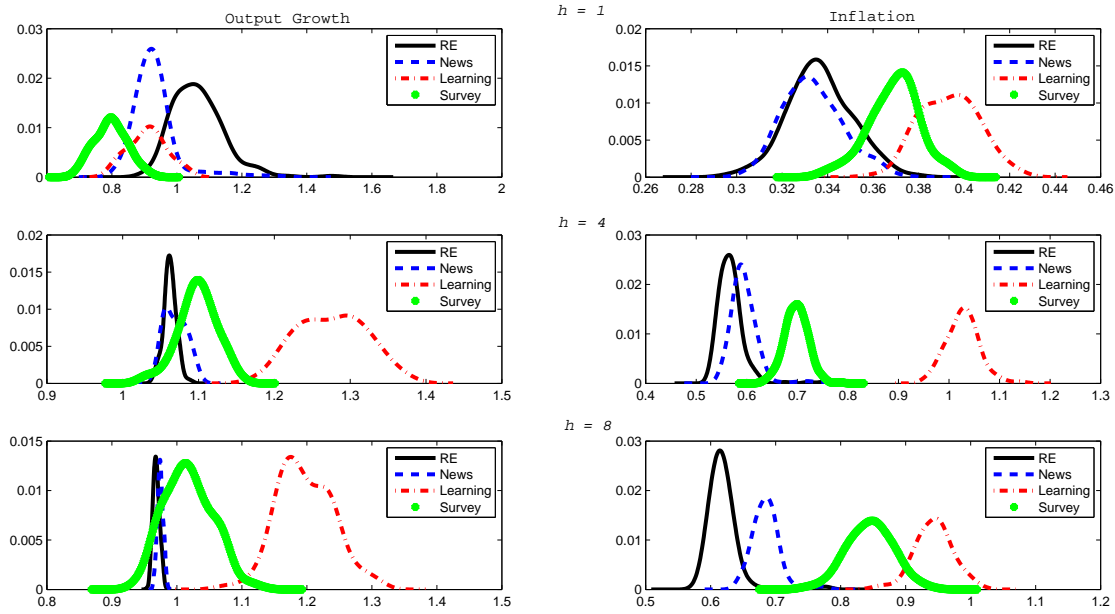


FIGURE 7. Univariate forecasting accuracy. *Note:* the figure shows the posterior distribution of the RMSE for output growth and inflation forecasts. The top panels refer to one-quarter-ahead forecasts, the middle panels to four-quarter-ahead forecasts, the bottom panels to eight-quarter-ahead forecasts. Each panel displays the posterior distributions obtained for the model under different expectation formation schemes.

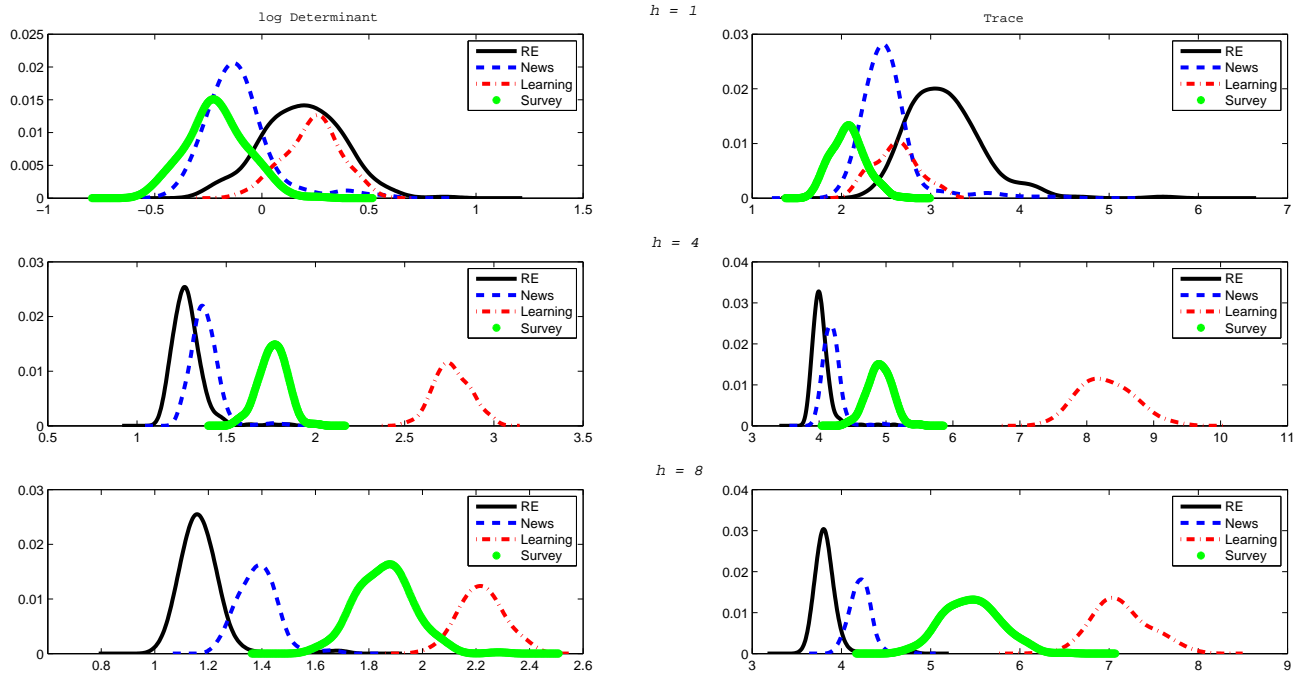


FIGURE 8. Multivariate forecasting accuracy. *Note:* the figure shows the posterior distributions of the log determinant (left) and of the trace of the forecast errors MSE matrix. The top panels refer to one-quarter-ahead forecasts, the middle panels to four-quarter-ahead forecasts, the bottom panels to eight-quarter-ahead forecasts. Each panel displays the posterior distributions obtained for the model under different expectation formation schemes.