

Expectations, Learning and Business Cycle Fluctuations

By STEFANO EUSEPI AND BRUCE PRESTON*

*This paper develops a theory of expectations-driven business cycles based on learning. Agents have incomplete knowledge about how market prices are determined and shifts in expectations of future prices affect dynamics. Learning breaks the tight link between fundamentals and equilibrium prices, inducing periods of erroneous optimism or pessimism about future returns to capital and wages which are partially validated by subsequent data. In a real business cycle model, the theoretical framework amplifies and propagates technology shocks. Moreover, it produces agents' forecast errors that are consistent with business cycle properties of forecast errors for a wide range of variables from the Survey of Professional Forecasters. JEL Classifications: E32, D83, D84
Keywords: Expectations-Driven Business Cycles, Learning Dynamics, Pigou Cycles*

Recently there has been renewed interest in shifting expectations as a source of business cycle fluctuation. A range of models have been explored that rely variously on multiple equilibria, exogenous news about future productivity and imperfect information — see, for example, Jess Benhabib and Roger Farmer (1994), Paul Beaudry and Frank Portier (2007), Nir Jaimovich and Sergio T. Rebelo (2008) and Guido Lorenzoni (2008). These frameworks seek not only to explain business cycle fluctuations with changes in expectations but also to resolve comovement problems that arise in real business cycle theory.

This paper explores an alternative theory based on learning dynamics. In the context of an otherwise standard stochastic growth model, we consider an environment in which households and firms have an incomplete model of the macroeconomy, knowing only their own objectives, constraints and beliefs. Agents are optimizing, have a completely specified belief system but do not know the equilibrium mapping between primitive disturbances, the aggregate state of the economy and market clearing prices. By extrapolating from historical patterns in observed data they approximate this mapping to forecast variables that are exogenous to their decision problems, such as the rental rate of capital and

* Eusepi: Macroeconomic and Monetary Studies Function, Federal Reserve Bank of New York, 33 Liberty St, New York NY 10045, stefano.eusepi@ny.frb.org. Preston: Department of Economics, Columbia University, New York NY 10027, bp2121@columbia.edu and Center for Applied Macroeconomic Analysis, Australian National University, Sir Roland Wilson Building #120, Canberra ACT 2601. The authors thank seminar participants at the The University of Auckland, Workshop in Macroeconomic Dynamics 2008, Australian National University, Central Bank of Brazil's conference "XI Annual Inflation Targeting Seminar", CAMA-ANU conference "Behavioural Macroeconomics: Theory and Implications", Cambridge University conference on "Learning and Macroeconomic Policy", Columbia University, ES-SIM 2008, Federal Reserve Bank of Atlanta, Hong Kong University, The Hong Kong University of Science and Technology, Northwestern University, NYU Stern, Reserve Bank of Australia's conference on "Monetary Policy in Open Economies", Society of Economic Dynamics 2007, Society of Computational Economics 2008, Stanford University, Universitat Autònoma and CREI-Universitat Pompeu Fabra. We thank Carlos Carvalho, Chris Edmond, Marc Giannoni, Chryssi Giannitsarou, Albert Marcet, Kristoffer Nimark, Aysegül Sahin, Jaume Ventura, two anonymous referees and particularly Florin Bilbiie for useful conversations and comments. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. The usual caveat applies.

the real wage. This belief structure has the property that beliefs affect the true data generating process of the economy which in turn affects belief formation. The economy is therefore self-referential: shifts in beliefs about future returns to labor and capital affect current market clearing prices which in turn can reinforce beliefs. In this environment, current prices can become less informative about future economic conditions generating fluctuations in real activity.

This kind of mechanism driving business cycle fluctuations is found in early writings on macroeconomic dynamics. For example, Alfred C. Pigou (1927), on page 122, writes:

"[...] a rise in prices, however brought about, by creating some actual and some counterfeit prosperity for business man, is liable to promote an error of optimism, and a fall in prices an error of pessimism, and this mutual stimulation of errors and price movements may continue in a vicious spiral until it is checked by some interference from outside."

Shifts in expectations, whether in part due to changes in fundamentals or in part due to error, are a source of business cycle fluctuation that may be self-fulfilling. We introduce learning into a canonical real business cycle model to generate expectations dynamics of this kind.¹ Learning breaks the tight link between fundamentals and, through expectations formation, equilibrium prices and allocations. This model property engenders periods of erroneous optimism or pessimism about future returns to capital and wages which are in part validated by subsequent data. An appealing feature is that agents' forecast errors of real wage and interest rates display similar business cycle properties as forecast errors for a wide range of variables from the Survey of Professional Forecasters. For example, in both the model and data, market participants under-predict interest rates during expansions and contrariwise during contractions.² A rational expectations version of the model is unable to explain such data. These mechanisms have the added benefit of generating additional amplification and propagation of technology shocks relative to a rational expectations version of the model.

The model is calibrated to match properties of post-war US quarterly data. One novel feature of the calibration is the use of survey data on forecasts to discipline the learning mechanism. Learning introduces one new parameter called the gain, which is restricted to a set of values such that the model replicates business cycle patterns observed in survey data on expectations. Given this restriction, the learning model in general provides a superior characterization of second-order moments of observed data than does the model under rational expectations. Additional specific results on amplification and propagation are as follows. First, learning amplifies technology shocks. Relative to a rational expectations analysis of the model, a 10 - 20 percent smaller standard deviation of technology shocks is required to match the standard deviation of HP-filtered output data.³ Moreover, the volatility of investment and hours relative to output is roughly 10 - 25 and 20 - 40 percent greater than under rational expectations. Second, the persistence properties of our model bear much closer resemblance to observed data. The first-order autocorrelation properties of output, hours and investment growth are well matched despite shocks being identically and independently distributed over time — hump-shaped impulse responses are observed. These features of the data are typically problematic for real business cycle theory as documented by Timothy Cogley and James Nason (1993) and Julio J. Rotemberg and Michael Woodford (1996).

¹The model will not have the property of "vicious spirals".

²This is also consistent with Monika Piazzesi and Martin Schneider (2008) which studies the role of subjective beliefs in bond pricing.

³Amplification is monotonic in the gain, and the range corresponds to the minimal and maximal values considered for this parameter.

The improvement in fit can be traced to shifting beliefs amplifying intertemporal substitution of consumption and leisure. The only source of exogenous variation are technology shocks, which have two effects. First, as in standard real business cycle theory, a technology improvement shifts the production frontier with well-understood implications. Second, in subsequent periods, households revise their beliefs in response to changed market opportunities. In particular, *relative* to rational expectations, households are more optimistic about the future path of returns to capital and more pessimistic about future returns to labor. The former leads to substitution of current for future consumption and a high marginal utility of income, an effect reinforced by the latter. Combined, these expectations effects induce a smaller increase in consumption, and, consequently, a larger shift in labor supply and investment in the period after the shock. This amplification of substitution effects in response to a technology shock relative to rational expectations explains the increased volatility in these variables. The delayed adjustment in beliefs explains the persistence.

Earlier analyses have explored learning as a source of amplification and propagation. Noah Williams (2003) examines a standard real business cycle model, concluding that adaptive learning is unlikely to help improve the performance of such models. Reproducing that analysis in the context of our model shows that this is indeed the case. The difference in conclusions stems from the failure to model optimal decisions conditional on maintained beliefs as done in Albert Marcet and Thomas J. Sargent (1989) and Bruce Preston (2005). Our paper relates to other recent contributions in the learning literature by Fabio Milani (2007), Eva Carceles-Poveda and Chryssi Giannitsarou (2007) and Kevin Huang and Zheng Liu and Tao Zha (2008). These papers are discussed later, but like Williams (2003), they consider models in which only one-period-ahead forecasts matter for household and firm behavior — decisions are not optimal given the maintained beliefs.⁴

This paper belongs to a long literature reconciling the predictions of real business cycle theory with observed data — see, *inter alia*, Gary Hansen (1985), Richard Rogerson (1988), Lawrence Christiano and Martin Eichenbaum (1992), Jess Benhabib and Roger Farmer (1994), David Andolfatto (1996), Craig Burnside and Martin Eichenbaum (1996), Charles T. Carlstrom and Timothy S. Fuerst (1997) and Stephanie Schmitt-Grohe (2000). These papers introduce various frictions, including indeterminacy of rational expectations equilibrium, to benchmark theory to improve the amplification and propagation of technology shocks. Our paper furthers this line of inquiry by considering learning dynamics as an alternative friction.

The introduction of imperfect information and learning in the real business cycle framework dates back to Finn E. Kydland and Edward C. Prescott (1982). In their model, the stochastic process for technology is composed of two unobserved shocks with different persistence. Agents face a signal extraction problem when predicting future productivity. More recently, Rochelle Edge and Thomas Laubach and John Williams (2007) show in a similar model that learning has important effects in the response of the economy to a persistent shift in productivity growth. In these models learning is not an endogenous source of propagation because changes in endogenous variables do not affect the agents' learning process. On the contrary, gradual recognition of the productivity changes generates a gradual response to the shock — a property determined by the specified signal-to-noise ratio in the exogenous process.

⁴These papers are similarly pessimistic about the potential of learning dynamics to resolve empirical questions. The exception is Milani (2007) which makes some progress in a New Keynesian model. Differing conclusions likely emerge from the role of initial conditions which do not play a role in our analysis.

I. A Simple Model

The following section details a stochastic growth model similar in spirit to Kydland and Prescott (1982), Edward C. Prescott (1986) and Robert G. King and Charles I. Plosser and Sergio T. Rebelo (1988). The major difference to this earlier literature is the incorporation of near-rational beliefs, delivering an anticipated utility model of the kind discussed by David Kreps (1998) and Thomas I Sargent (1999). The analysis follows Albert Marcet and Thomas J. Sargent (1989) and Bruce Preston (2005), solving for optimal decisions conditional on current beliefs.

A. Microfoundations

Firms. There is a continuum of identical competitive firms of mass one. Each firm i produces the economy's only good using capital K_t^i and labor H_t^i as inputs according to the production function

$$(1) \quad Y_t^i = (K_t^i)^\alpha (X_t H_t^i)^{1-\alpha}$$

where $0 < \alpha < 1$. The stochastic trend X_t is labor-augmenting technical progress, identical for each firm, evolving according to $\ln(X_{t+1}/X_t) = \gamma_t = \ln \bar{\gamma} + a_{t+1}$ where a_t is an independent, identically distributed random variable with zero mean and standard deviation σ_A and $\bar{\gamma} > 0$. Each firm chooses labor and capital inputs to maximize profits $\Pi_t^i = Y_t^i - R_t^K K_t^i - W_t H_t^i$ taking factor prices as given. The first-order conditions to a firm's optimization problem provide

$$(2) \quad W_t = (1 - \alpha) (K_t^i)^\alpha (X_t)^{1-\alpha} (H_t^i)^{-\alpha}$$

$$(3) \quad R_t^K = \alpha (K_t^i)^{\alpha-1} (X_t H_t^i)^{1-\alpha}$$

which equate factor prices with their real marginal products. R_t^K is the rental rate of capital and W_t is the real wage.

Households. Households maximize their intertemporal utility derived from consumption and leisure

$$(4) \quad \hat{E}_t^j \sum_{T=t}^{\infty} \beta^{T-t} \left[\ln C_T^j - \nu(H_T^j) \right]$$

subject to the flow budget constraint

$$(5) \quad C_t^j + K_{t+1}^j = R_t^K K_t^j + W_t H_t^j + (1 - \delta) K_t^j$$

where C_t^j denotes household j 's consumption, K_t^j the holdings of the aggregate capital stock available at the beginning of period t , with $K_0^j > 0$ given; and H_t^j represents the fraction of the available time (normalized to one unit per period) spent on non-leisure activities. The function $\nu(\cdot)$ is convex. The functional forms are chosen to be consistent with a balanced growth path. Households supply labor and capital in perfectly competitive markets. The household's discount factor and the depreciation rate of capital satisfy $0 < \beta, \delta < 1$.

The expectation operator \hat{E}_t^j denotes agent j 's subjective beliefs. In forming expectations, households and firms observe only their own objectives, constraints and realizations

of aggregate variables that are exogenous to their decision problems and beyond their control. The agent's problem is to choose sequences of consumption, hours worked, and capital in order to maximize (4) subject to (5), taking as given prices and the capital stock available at the beginning of the period. Beliefs are specified in the next section.

Household optimization yields the well-known conditions for expected consumption growth and hours worked

$$W_t = C_t^j \nu_H \left(H_t^j \right) \quad \text{and} \quad \hat{E}_t^j \left[\beta \frac{C_t^j}{C_{t+1}^j} (R_{t+1}^K + (1 - \delta)) \right] = 1.$$

The paper's primary goal is the quantitative evaluation of the model. Following Finn E. Kydland and Edward C. Prescott (1982), employ a log-linear approximation of the model around a balanced growth path. For any variable G_t define $g_t = G_t/X_t$ as the corresponding normalized variable. Households observe the stochastic trend X_t and its distribution. This assumption is necessary to obtain stationary decision rules. However, households do not know what determines the evolution of X_t : that is, firms' technology. This is further discussed in section 3 of the paper. Balanced growth requires consumption, investment, output, the capital stock and real wages to grow at the rate of the stochastic trend so that

$$y_t = \frac{Y_t}{X_t}; \quad c_t = \frac{C_t}{X_t}; \quad i_t = \frac{I_t}{X_t}; \quad w_t = \frac{W_t}{X_t} \quad \text{and} \quad k_t = \frac{K_t}{X_{t-1}}$$

are stationary. Hours and the rental rate of capital are stationary on the balanced growth path. Details of the steady state and the log-linear approximation are confined to the online appendix.

Log-linearizing, solving the flow budget constraint forward, imposing the transversality condition and substituting for hours gives the intertemporal budget constraint

$$\epsilon_c \hat{E}_t^j \sum_{T=t}^{\infty} \beta^{T-t} c_T^j = \beta^{-1} \hat{k}_t^j + \hat{E}_t^j \sum_{T=t}^{\infty} \beta^{T-t} \left[\epsilon_w \hat{w}_T + \bar{R} \hat{R}_T^K - \beta^{-1} \hat{\gamma}_T \right].$$

The coefficients ϵ_c and ϵ_w are constants that are composites of model primitives and $\bar{R} > 0$. This relation states the expected present value of consumption must be equal to the capital stock available at the beginning of the period plus the expected present value of wage and rental income. These latter variables are outside the control of the household, given the assumption of competitive markets.

To determine optimal consumption decisions, combine the intertemporal budget constraint with a log-linear approximation to the consumption Euler equation to yield

$$(6) \quad \begin{aligned} \hat{c}_t^j &= \frac{1 - \beta}{\epsilon_c} \left[\beta^{-1} \left(\hat{k}_t^j - \hat{\gamma}_t \right) + \bar{R} \hat{R}_t^K + \epsilon_w \hat{w}_t \right] \\ &+ \hat{E}_t^j \sum_{T=t}^{\infty} \beta^{T-t} \left[\frac{(1 - \beta)}{\epsilon_c} - \beta \right] \beta \bar{R} \hat{R}_{T+1}^K + \hat{E}_t^j \sum_{T=t}^{\infty} \beta^{T-t} \frac{(1 - \beta)}{\epsilon_c} \beta \epsilon_w \hat{w}_{T+1}. \end{aligned}$$

The consumption decision rule comprises three terms. The first shows the impact that the current level of the capital stock and current prices have on consumption. The second and third terms show how expected variations in permanent income affect current consumption. The former has two parts corresponding to the positive income effect and the negative substitution effect of higher returns to capital on current consumption. The latter has only one part as the income and substitution effects of a wage increase both

raise current consumption.

B. Market clearing and aggregate dynamics

As households have the same preferences and constraints; firms the same technology; and beliefs are assumed homogeneous across all agents (although they are assumed not to be aware of that) the analysis considers a symmetric equilibrium in which $\hat{k}_t^i = \hat{k}_t^j = \hat{k}_t$; $\hat{H}_t^j = \hat{H}_t^i = \hat{H}_t$; $\hat{v}_t^i = \hat{v}_t^j = \hat{v}_t$ for all i, j, t . Integrating over the continuum provides aggregate consumption demand

$$(7) \quad \begin{aligned} \hat{c}_t &= \frac{1-\beta}{\epsilon_c} \left[\beta^{-1} \hat{k}_t + \bar{R} \hat{R}_t^K - \beta^{-1} \hat{\gamma}_t + \epsilon_w \hat{w}_t \right] \\ &+ \hat{E}_t \sum_{T=t}^{\infty} \tilde{\beta}^{T-t} \left[\frac{(1-\beta)}{\epsilon_c} - \beta \right] \beta \bar{R} \hat{R}_{T+1}^K + \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} \frac{(1-\beta)}{\epsilon_c} \beta \epsilon_w \hat{w}_{T+1} \end{aligned}$$

where $\int \hat{E}_t^j dj = \hat{E}_t$ denotes average expectations in the population. Aggregate consumption dynamics inherit the properties of individual decision rules. This is the only model equation that depends on expectations, and therefore of central focus. If near-rational expectations are to be a source of amplification and propagation of shocks, the effects must originate here.

A log-linear approximation to relations (1), (2), (3), (5) and labor-leisure condition yields the remaining model equations.

II. Beliefs

Optimal decisions require households to forecast the evolution of future wages and returns to capital. They are assumed to use a simple econometric model, relating wages in efficiency units and the capital rental rate to the aggregate stock of capital also expressed in efficiency units. That is

$$(8) \quad \hat{R}_t^K = \omega_0^r + \omega_1^r \hat{k}_t + e_t^r,$$

$$(9) \quad \hat{w}_t = \omega_0^w + \omega_1^w \hat{k}_t + e_t^w$$

and

$$(10) \quad \hat{k}_{t+1} = \omega_0^k + \omega_1^k \hat{k}_t + e_t^k$$

where e_t^r , e_t^r and e_t^r are regression errors. The beliefs contain the same variables that appear in the minimum-state-variable rational expectations solution to the model. And, while the rational expectations solution does not contain a constant, it has the natural interpretation under learning of capturing uncertainty about the steady state.

Rational Expectations. The model solution under rational expectations implies (to a first-order approximation) that labor and capital prices and the next-period capital stock expressed in efficiency units are linearly related to aggregate capital, with *time-invariant* coefficients $\omega_0^r = \omega_0^w = \omega_0^k = 0$ and $\omega_1^r = \bar{\omega}_1^r$, $\omega_1^w = \bar{\omega}_1^w$, $\omega_1^k = \bar{\omega}_1^k$. The agents' forecasting model nests beliefs that would be observed in a rational expectations equilibrium. Under rational expectations $e_t^r = \bar{\omega}_3^r \hat{\gamma}_t$, $e_t^w = \bar{\omega}_3^w \hat{\gamma}_t$ and $e_t^k = \bar{\omega}_3^k \hat{\gamma}_t$.

Perpetual learning. Agents estimate equations (8) – (10), updating their coefficient estimates every period as new data become available. Following recent literature, households update their estimates using a discounted least-squares estimator, assigning lower weight to older observations to protect against structural change.⁵ Letting $\omega' = (\omega_0, \omega_1)$, $z_t = \left(\hat{R}_t^K, \hat{w}_t, \hat{k}_{t+1} \right)$ and $q_{t-1} = \left(1, \hat{k}_t \right)$, the algorithm can be written in recursive terms as

$$(11) \quad \hat{\omega}_t = \hat{\omega}_{t-1} + gR_t^{-1}q_{t-1} \left(z_t - \hat{\omega}'_{t-1}q_{t-1} \right)'$$

$$(12) \quad R_t = R_{t-1} + g \left(q_{t-1}q'_{t-1} - R_{t-1} \right)$$

where $\hat{\omega}_t$ denotes the current-period's coefficient estimate and $g \in (0, 1)$ denotes the constant gain, determining the rate at which older observations are discounted. The constant gain assumption delivers perpetual learning, as market participants 'forget' the past.

This learning rule is motivated by the assumption that agents perceive factor prices and capital to be non-stationary. Observed changes in prices are in part attributed to shifts in the model parameters ω_t and in part to the idiosyncratic disturbances e_t^r , e_t^w and e_t^k . Therefore, any change in factor prices and capital is partially interpreted as a permanent change in the long-run equilibrium values of the variables they forecast. The constant gain g regulates how much weight agents attribute to these permanent changes. The model nests rational expectations: as $g \rightarrow 0$, agents perceive all changes as temporary and their model converges to the true data generating process of the economy.⁶ One advantage of this simple model is that departure from rational expectations is regulated by a single parameter. In the on-line appendix we offer a Kalman filter interpretation to the simple econometric model discussed here.

Information. It is assumed the innovation, $\hat{\gamma}_t$, is not used in equations (8) – (10). This does not imply $\hat{\gamma}_t$ is unobserved — indeed, (6) implies consumption decisions are in part determined by these innovations. If the innovation was used in forecasting, agents would not face an inference problem and learn quickly given that the only disturbance in the model is the technology shock.⁷ Not including this disturbance in the econometric model is justifiable because, while individual households and firms observe these disturbances, they do not how they are mapped into market clearing prices in general equilibrium. This is a natural implication of the maintained assumptions. Agents do not know individual preferences and technologies of other market participants. Households understand that there is a stochastic trend but know neither its determinants nor that firms are identical. Wherefore, they cannot infer either an aggregate production function or the precise relationship between factor prices, aggregate capital and technology shocks.

⁵Of course we consider an otherwise stationary model environment with a single shock so as to clearly isolate the role of expectations in generating business cycle fluctuations. Adding structural change would generate further volatility.

⁶The convergence properties of this class of model are discussed in George W. Evans and Seppo Honkapohja (2001).

⁷Formally, including the disturbance would generate a singularity in the regression if initial beliefs coincide with the rational expectations equilibrium. When initial beliefs differ from the rational expectations equilibrium, the regression is well defined, but because there is no uncertainty about the forecasting model, beliefs quickly converge to the predictions of a rational expectations analysis where the singularity would again emerge given infinite data. And with a small gain, as in our analysis, the regression's variance-covariance matrix would still be close to singular.

Suppose the process X_t is in fact determined by

$$X_t = \prod_{i=1}^N Z_t^i$$

where each Z_t^i is random walk with drift unobserved by agents. These individual processes possibly reflect population growth, technology growth in other sectors of the economy, or long-term shifts in the aggregate labor supply. Because agents can not distinguish primitive processes they can not estimate each disturbance's contribution to the determination of prices. Under rational expectations, the equilibrium return to capital would be determined by

$$\hat{R}_t^K = \bar{\omega}_0^* + \bar{\omega}_1^* \hat{k}_t + \sum_{i=1}^N \bar{\omega}_{1+i}^* \hat{\gamma}_t^{z,i}$$

with obvious notation. Our assumption is a reduced-form representation of a more complicated model in which agents make inferences in the presence of multiple imperfectly observed disturbances.

Recourse to what some may view as a stark assumption is solely due to the lack of meaningful heterogeneity in the model. Indeed, this assumption is similar to the imperfect common knowledge literature where heterogeneity in information is explicitly modelled but where it is often assumed that only certain kinds of aggregate data are public knowledge or only certain markets are available to trade state-contingent claims. Absent these assumptions prices would fully reveal information about which agents are assumed to have only imperfect understanding — there is no inference problem. For example, see Lorenzoni (2008).

Timing. In any period t agents inherit belief parameters determined by period $t - 1$ data. While the forecast function is predetermined, expectations themselves are not. Agents observe the same variables that a ‘rational’ agent would observe. The only difference is that their learning shifts the forecast function over time. For example, the one-period-ahead forecast of \hat{R}_t^K is

$$\hat{E}_t \hat{R}_{t+1}^K = \hat{\omega}_{0,t-1}^r + \hat{\omega}_{1,t-1}^r \hat{k}_{t+1}$$

where $\hat{\omega}_{0,t-1}^r$ and $\hat{\omega}_{1,t-1}^r$ are the previous-period's estimates of belief parameters that define the period t forecast function. As \hat{k}_{t+1} is predetermined in period t it is observed, while future levels of the capital stock must be forecast. Finally in forecasting over the decision horizon agents do not take into account that they update their beliefs in subsequent periods. The model is one of anticipated utility — see Sargent (1999).

True Data Generating Process. Using (8) – (10) to substitute for expectations in (7) and solving delivers the actual data generating process

$$(13) \quad z_t = T_1(\hat{\omega}_{t-1}) q_{t-1} + T_2(\hat{\omega}_{t-1}) \hat{\gamma}_t$$

$$(14) \quad \hat{\omega}_t = \hat{\omega}_{t-1} + g R_t^{-1} q_{t-1} \left(\left[T_1(\hat{\omega}_{t-1}) - \hat{\omega}'_{t-1} \right] q_{t-1} + T_2(\hat{\omega}_{t-1}) \hat{\gamma}_t \right)'$$

$$(15) \quad R_t = R_{t-1} + g (q_{t-1} q'_{t-1} - R_{t-1})$$

and

$$(16) \quad \begin{bmatrix} \hat{c}_t & \hat{i}_t & \hat{H}_t \end{bmatrix}' = \Psi z_t,$$

where $T_1(\hat{\omega})$ and $T_2(\hat{\omega})$ are nonlinear functions of the previous-period's estimates of beliefs and Ψ is a matrix comprised of composites of primitive model parameters. The actual evolution of z_t is determined by a time-varying coefficient equation in the state variables \hat{k}_t and $\hat{\gamma}_t$, where the coefficients evolve according to (14) and (15). The model given by relations (13) - (16) has a unique non-stochastic steady state corresponding to the steady state under rational expectations.⁸ George W. Evans and Seppo Honkapohja (2001) show that for a gain sufficiently close to zero the distribution of the estimates $\hat{\omega}_t$ is normal and centered around the time-invariant coefficients of the rational expectations equilibrium.

Forecast Errors. The evolution of z_t depends on $\hat{\omega}_{t-1}$, while at the same time $\hat{\omega}_t$ depends on z_t . Learning induces self-referential behavior. Comparing (8) - (10) and (13) - (16) it is immediate that market participants have a mis-specified model of the economy: that is $\omega \neq T_1(\omega)$. In particular, the economy exhibits time variation in the coefficients $T_1(\hat{\omega}_{t-1})$ and $T_2(\hat{\omega}_{t-1})$ but their evolution differs to the agents' model. The mis-specification stems from the fact that individuals fail to internalize the impact of updating their model on the aggregate economy. Each market participant views structural change as exogenous, ignorant that every other agent is using the same updating rule.⁹ In fact, the time variation in the coefficients $T_1(\hat{\omega}_{t-1})$ and $T_2(\hat{\omega}_{t-1})$ is endogenously generated by the agents' learning process. Model mis-specification is the source of the systematic forecast errors discussed next.

III. Calibration

The sample characteristics we seek to match are for US data, 1948:Q1 to 2007:Q4. A short description of each series is contained in the Appendix. We set the discount rate to $\beta = 0.99$. We assume separable preferences between consumption and leisure, with log-utility for consumption and close-to-linear disutility of labor.¹⁰ Firms' technology is specified by a capital share $\alpha = 0.34$ and steady-state growth rate of labor augmenting technical progress equal to $\bar{\gamma} = 1.0053$, consistent with the quarterly mean output growth over the sample.

Two parameters are left to calibrate: the standard deviation of the shock, σ_A , and the constant gain, g . The former is calibrated by minimizing the sum-of-squared distances between the model-implied volatility of HP-detrended output and the corresponding data moment.¹¹ The latter parameter requires further discussion. Households require forecasts of future wages and rental rates to make current spending decisions. The gain parameter regulates the properties of these expectations. Ideally it would be disciplined by data on these forecasts. Unfortunately, mapping the model concepts of wage and interest rates to measured data concepts is problematic. There are scarce survey forecast data on wages and certainly no data on wages normalized by the level of technology, the object that is actually forecast by households in our model. And while survey expectations of nominal interest rates at various maturities and inflation permit inferences about forecasted real

⁸Although the agents' model assumes a non-stationary environment, the actual evolution of the economy is a stationary process with time-varying parameters.

⁹In the on-line appendix we show that according to the Kalman filter interpretation of agents updating algorithm, the coefficients ω_t are assumed to be evolving according to a random walk.

¹⁰Accordingly, the inverse Frisch elasticity of labor supply satisfies $\epsilon_H \rightarrow 0$. This approximates the labor supply properties of a model of indivisible labor — see Hanson (1985) and Rogerson (1988).

¹¹The time series are generated by simulating the model for $2000 + T$ periods. The first 2000 periods guarantee convergence to the model stationary distribution and are discarded. The simulation is repeated 5000 times. The time of the simulation is set to $T = 160$ quarters which corresponds to the sample size for US data.

interest rates, they are not necessarily directly related to the marginal product of capital — see Robert E. Hall (1988) and Mulligan (2002).

Given these limitations, the choice of gain parameter is disciplined indirectly using survey data on expectations. Table 1, Panel A, documents properties of the forecast errors from the Survey of Professional Forecasters for measures of economic activity, goods prices and nominal interest rates. The data are described in the Appendix. Columns 1 and 3 give the autocorrelation of forecasts errors at a quarterly and annual frequency, while columns 2 and 4 give the predictable movements in these same errors conditional on lagged real GDP growth.¹²

Table 1 Forecast Properties

	Statistic			
	$(FE_t^{1Q}, FE_{t-1}^{1Q})$	$(FE_t^{1Q}, \Delta \mathbf{y}_{t-1})$	$(FE_t^{4Q}, FE_{t-4}^{4Q})$	$(FE_t^{4Q}, \Delta \mathbf{y}_{t-4})$
Panel A: Survey of Professional Forecast Data				
Economic activity				
Real. RGDP ¹	0.11 – 0.26	–	–	–
Nominal GDP	0.14	–	0.1	–
Real GDP ²	0.37	–	0.18	–
Industrial prod.	0.44	–	0.10	–
Unemp. rate	0.58	0.38	0.14	0.37
Interest rates				
Tbill	0.27	–0.17	0.10	–0.43
Tbond (10yrs)	0.32	–0.32	–0.45	–0.20
Corporate bond	0.40	–0.20	–0.21	–0.22
Real (ex post)	0.38	–0.05	–	–
Prices (inflation)				
GDP Deflator	0.56	0.04	0.53	–0.2
CPI	0.14	–0.24	–0.08	–0.27
Panel B: Model Predictions				
Learning				
Sample: 162 Q				
R_t^k forecast	0.36 (0.19)	–0.09 (0.11)	+0.03 (0.15)	–0.02 (0.09)
w_t forecast	0.36 (0.19)	0.09 (0.11)	–0.03 (0.15)	0.02 (0.09)
Sample: 110 Q				
R_t^k forecast	0.31 (0.21)	–0.09 (0.12)	+0.01 (0.17)	0.00 (0.11)
w_t forecast	0.31 (0.21)	0.09 (0.12)	–0.01 (0.17)	0.00 (0.11)
REE				
Sample: 110 Q				
R_t^k forecast	–0.01 (0.08)	0.00 (0.09)	+0.03 (0.13)	–0.02 (0.09)
w_t forecast	–0.01 (0.08)	0.00 (0.09)	–0.03 (0.13)	0.02 (0.09)

Note: data from Survey of Professional Forecasters. We use the median one- and four-quarters ahead forecasts. The variable Real GDP¹ is obtained as the difference between the forecast of nominal GDP

¹²Forecast errors for any variable Z_t are defined as $FE_t^{1Q} = E_t Z_{t+1} - Z_{t+1}$ for one-quarter-ahead forecasts and $FE_t^{4Q} = E_t Z_{t+4} - Z_{t+4}$ for four-quarter-ahead forecasts.

and the forecast of the GDP deflator. For robustness, we computed the one-quarter-ahead forecast errors using alternative data vintages available to forecasters. The autocorrelation of one-period-ahead forecast errors is 0.26 for the first vintage, 0.16 for the second vintage, 0.22 for the fifth vintage, 0.19 for the ninth vintage and 0.12 for the latest available vintage. The variable Real GDP² is the forecast for real GDP from the survey, which is available for a smaller sample period. The forecast for the ex-post real rate is obtained by subtracting the two-periods-ahead forecast for the GDP deflator from the one-period-ahead forecast of the 3-months Tbill.

At a quarterly frequency, forecast errors are positively autocorrelated, a pattern for the most part reflected in the data at annual frequencies — exceptions being for two forecast measures of interest rates which display negative autocorrelation and the CPI.¹³ Forecast errors also show a systematic pattern over the business cycle. For example, market participants under-predict interest rates and over-predict unemployment during expansions and contrariwise during contractions. Counter cyclicity in interest rate forecast errors has also been documented by Monika Piazzesi and Martin Schneider (2008).

We use the above evidence to impose discipline on the choice of gain. The benchmark model assumes a gain of 0.002. While the discussion gives focus to this particular value, later robustness exercises gauge the sensitivity of this calibration. It is shown that a range of gain parameters engender learning dynamics that have the property that their implied forecast errors are consistent with both the above evidence and amplification and propagation of technology shocks. In this sense, the Survey of Professional Forecast data restricts the gain parameter to a set of values.

Our calibrated gain is considerably smaller than values found in the literature, which range from 0.007 – 0.05 — see, for example, William A. Branch and George W. Evans (2006) and Fabio Milani (2007), which estimate the gain, and Athanasios Orphanides and John C. Williams (2005). Finally, to interpret this magnitude, note the gain indexes the weight assigned to past data. This value of the gain implies that observations that are 50 years old receive a weight of $(1 - 0.002)^{200} \simeq 0.67$, implying agents do not discount past data too heavily.¹⁴

IV. Central Results

A. Inspecting the Mechanism: the Effects of a Technology Shock

To develop intuition on the role of learning in business cycle fluctuations, consider the model's impulse response functions to a unit technology shock. For stationary variables, the impulse response functions are expressed in percentage deviations from steady state. For non-stationary series, the impulse responses are reported in percentage deviations from the balanced growth path. We plot the dynamics of these series relative to their paths in absence of the shock. A unit positive technology shock leads to a unit increase in the level of these series. In the model with learning the effects of a disturbance depend on the precise beliefs maintained by households at the time of the shock.¹⁵ In each

¹³Several papers have documented that survey forecasts as measured by surveys like Michigan, Livingston and SPF fail to be consistent the 'rational' expectations in terms of unbiasedness and serially uncorrelated forecast errors. Most studies focus on inflation expectations — see, for example, Lloyd B. Thomas (1999) and John M. Roberts (1997) and references therein.

¹⁴For this value of g , agents would give approximately zero weight to observations that are 500 years old.

¹⁵Impulse response functions for the learning model are generated by simulating the model twice for $2000 + T$ periods. The first 2000 periods guarantee convergence to the model stationary distribution and are discarded. The second simulation includes a unit shock in period 2001. The T -period impulse response to a unit technology shock is then given by the difference between these two trajectories. The simulation is repeated 5000 times. A projection facility is employed to ensure beliefs are not explosive.

plot the solid lines correspond to the median point-wise impulse response function, while the dotted lines provide the interquartile range of the simulated impulse responses. The dashed line gives the corresponding impulse response predicted by a rational expectations analysis of the model.

Figure 1 reports the impulse response functions for output, consumption, investment and hours. That learning amplifies disturbances relative to rational expectations is immediate. Output, hours and investment display a hump-shaped profile in response to a technology shock. The precise quantitative implications are documented below. An interesting feature of the model concerns dynamics in the period after the technology shock dissipates. In a rational expectations equilibrium, all model variables, appropriately normalized, are a linear function of the capital stock and the disturbance to the growth rate of technology. As the disturbance is assumed to be i.i.d., the observed dynamics one period after the shock are entirely determined by adjustment in the capital stock. Under learning, this is not the case. The technology shock leads to revisions in beliefs that commence the period after the disturbance. Subsequent dynamics are largely driven by revisions to beliefs. The following describes the sequential response of the economy in greater detail.

Impact For all series the impact effects of a technology shock are almost identical when comparing the median impulse response under learning and the impulse response under rational expectations. This is because agents' beliefs are distributed around the rational expectations prediction function, as shown in the next section. However, in the case of learning, there is variation in the impact effects. Depending on the precise beliefs of households and firms at the time of the shock, which along with the capital stock determine the state of the economy, the impact effect of the technology shock could be larger or smaller.

Individual forecasts In the period *after* the disturbance, agents revise upwards their beliefs about the returns to investment and downwards their beliefs about wages. A positive disturbance leads to an increase in the coefficients of (9) and a decrease in the coefficients of (8). The observed rise and fall of real interest rates and efficiency wages is in part attributed to the idiosyncratic shocks and in part to both changes in their long-run equilibrium returns and the elasticity of each rate with respect to capital. Consequently, the present discounted value of the returns to capital rise and the present discounted value of labor returns in efficiency units fall relative to the previous period, as indicated in Figure 2 which plots these infinite sums.¹⁶

Individual consumption and labor supply According to the optimal decision rule (7), optimism about future returns to capital tilts the consumption profile towards greater future consumption. This and the flatter expected wage path serve to increase the marginal utility of income, inducing intertemporal substitution of consumption and leisure. Consequently, individual investment and labor supply increase.

In practice this is invoked extremely rarely, reflecting the remarkable stability of the real business cycle model under learning and the maintained parameter assumption. This is also reflected in the ergodic distribution of beliefs in figure 4. Our results are not driven by explosive beliefs. As an additional check, we confirm that the mean and median of reported statistics are comparable.

¹⁶This does not imply that expected wages are lower than before the technology improvement. This can be inferred from the impulse response to consumption which under the current calibration is the same as the real wage. An analogous plot to figure 3 for efficiency wages reveals that the expected fall in efficiency wages is less than the realized fall, implying that the expected and actual level of wages are above the pre-shock level. See also the discussion of figure 3.

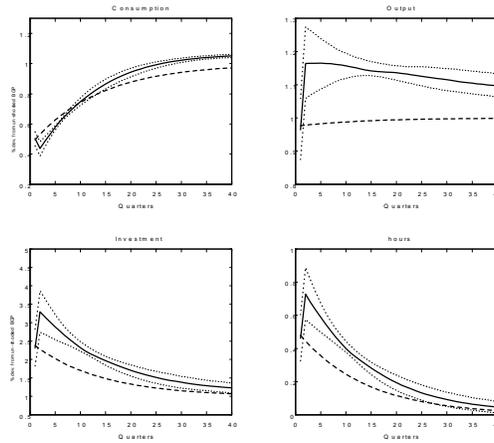


FIGURE 1. IMPULSE RESPONSES FOR CONSUMPTION, OUTPUT, INVESTMENT AND HOURS TO A 1% CHANGE IN TECHNOLOGY. THE SOLID LINE DENOTES THE MEDIAN IMPULSE RESPONSES FOR THE MODEL UNDER LEARNING. THE DASHED LINE SHOWS THE IMPULSE RESPONSE UNDER RATIONAL EXPECTATIONS. THE DOTTED LINE INDICATE THE 25TH AND 75TH PERCENTILE IMPULSE RESPONSE UNDER LEARNING.

Aggregate response The increase in aggregate labor supply lowers equilibrium wages, increasing the amount of hours worked and boosting the marginal productivity of capital. The realized return to capital turns out to be higher than agents’ forecasts, as shown in the top panel of Figure 3. Similarly, higher aggregate labor supply induces lower efficiency wages than expected. This “inertia” in the forecasts arises because individual agents fail to internalize the effects of their beliefs revision on the aggregate economy.¹⁷ However, agents’ predictions are in part realized in equilibrium outcomes in the period after the shock: the model generates dynamics that are broadly consistent with those described by Pigou (1927).

Convergence In subsequent periods agents revise their beliefs about the future evolution of the returns to capital and labor: estimated coefficients in (8) and (9) converge *monotonically* to their long-run values and the model converges to its unique steady state given no further shocks.¹⁸

Figure 3 presents a final set of impulse response functions that further elucidate the mechanics of the model. The top panel gives the dynamics of the difference between output under learning and rational expectations (dashed line) — revealing the additional amplification and propagation — and the model-implied forecast error of the rental rate. The forecast error for efficiency wages is the mirror image, with opposite sign. Together they underscore some important properties of the model. Even though households are optimistic about returns to capital *relative* to rational expectations, in equilibrium they

¹⁷This is captured by the fact that, under learning, $T_1(\omega) \neq \omega$ in (13).

¹⁸This monotonic convergence is related to the convergence properties of agents’ beliefs. The stability properties of the model imply real eigenvalues — see appendix.

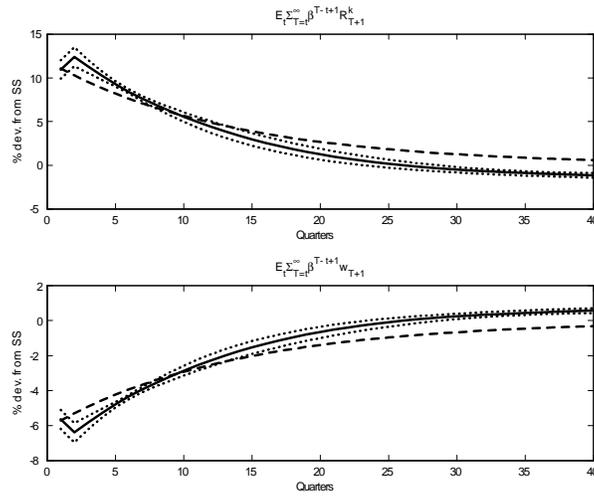


FIGURE 2. IMPULSE RESPONSES FOR FOR THE DISCOUNTED SUM OF R^k AND w TO A 1% CHANGE IN TECHNOLOGY. THE SOLID LINE DENOTES THE MEDIAN IMPULSE RESPONSES FOR THE MODEL UNDER LEARNING. THE DASHED LINE SHOWS THE IMPULSE RESPONSE UNDER RATIONAL EXPECTATIONS. THE DOTTED LINE INDICATE THE 25TH AND 75TH PERCENTILE IMPULSE RESPONSE UNDER LEARNING.

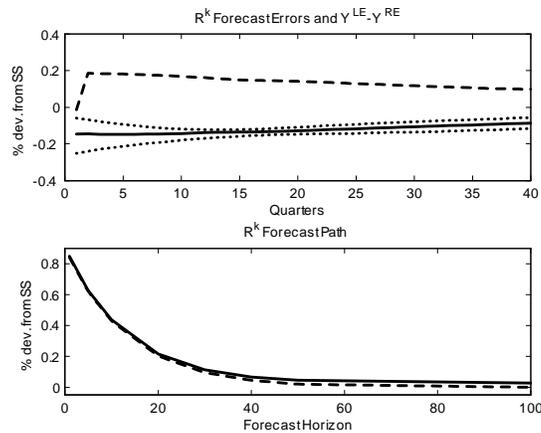


FIGURE 3. IMPULSE RESPONSE TO A 1% INCREASE IN TECHNOLOGY. THE TOP PANEL SHOW THE MEDIAN RESPONSE OF THE ONE-QUARTER AHEAD FORECAST ERROR FOR R^k (SOLID LINE) AND ITS INTERQUARTILE RANGE (DOTTED LINE) AND THE DIFFERENCE BETWEEN THE OUTPUT RESPONSE UNDER LEARNING AND UNDER RATIONAL EXPECTATIONS (DASHED LINE). THE BOTTOM PANEL SHOWS THE FORECASTED PATH OF R^k AT DIFFERENT FORECASTING HORIZONS, ON THE PERIOD AFTER THE TECHNOLOGY SHOCK. THE SOLID LINE REPRESENTS THE MODEL UNDER LEARNING AND THE DASHED LINE THE MODEL UNDER RATIONAL EXPECTATIONS.

still under-predict actual interest rates. The model displays a systematic relationship between forecast errors and the business cycle which is in line with the empirical evidence presented in the previous section. The bottom panel shows the forecasted path for R^k in the period after the shock, that is after agents' update the regression coefficients in equations (8) – (10). The forecast path under learning (solid line) and rational expectations (dashed line) are significantly different only at longer forecast horizons (more than 20 quarters). Under rational expectations agents' forecasting models have fixed coefficients while under learning the coefficients fluctuate over time. This important difference emerges at longer forecasting horizons because variation in the belief coefficients become the most important determinant of expectations (as model variables are expected to be close to steady state levels).

B. Statistical Properties

Model-implied forecast errors. Panel B, Table 1, shows forecast error statistics for each of three models: the learning model simulated for $T = 162$ quarters; the learning model simulated for $T = 110$ quarters; and the rational expectations model simulated for $T = 110$ quarters.¹⁹ These sample periods correspond to the two most frequent samples for which forecasts are available from the Survey of Professional Forecasters. The number reported in brackets is the standard deviation arising from sampling variability. Our baseline calibration implies serial correlation in agents' one-period-ahead forecast errors which is comparable to the Survey of Professional Forecast data reported in section 4. Moreover, the sampling variability easily encompasses these latter values. In contrast, the rational expectations model generates no persistence in forecast errors, even in small sample.

Several additional points are worthy of comment. First, longer model-simulated samples generate greater persistence in forecasts errors. This is because the low frequency movement in beliefs becomes a more dominant source of variation as the sample size increases. With a single technology shock, the serial correlation becomes highly persistent. Second, forecast errors of both wages and rental rates display correlation with past output growth. Again, the rational expectations model finds little predictable movement. Third, agents' four-quarter-ahead forecast errors display no pattern; at longer horizons model predictions are closer to rational expectations than to survey data.

Business cycle statistics. Table 2 reports summary statistics on the cyclical properties of various US data series and the model under both rational expectations and learning dynamics. The sample length of the simulations is $T = 162$. For each variable, Panel A reports the relative standard deviation and correlation with output for HP-filtered series (facilitating comparison to earlier studies based on filtered data). Panel B reports corresponding statistics for the growth rates of each series (which is more natural given the assumed stochastic trend). It also gives first-order serial correlation of growth rates along with the standard deviation arising from sampling variability in parenthesis.

Panel A shows learning dynamics amplify the effects of technology shocks. To match the variance of output, the learning model requires a technology disturbance with a standard deviation that is about 14 percent smaller than required under rational expectations. Moreover, the relative volatility of hours and investment is 16 and 29 percent higher respectively, bearing closer resemblance to data-implied moments than the rational expectations model. The former represents a significant success, being problematic for standard real business cycle theory — see Hansen (1985) and Rogerson (1988). In regards to consumption, wages and labor productivity, the model performs less well. Given the

¹⁹The model is simulated for 2000+T periods. The first 2000 observations are discarded and the statistics are computed using the remaining T observations. See footnote 16.

high elasticity of labor supply and the assumption of perfectly competitive markets the model predicts $\hat{C}_t \approx \hat{w}_t = \hat{Y}_t - \hat{H}_t$ and is therefore too stylized to capture the different dynamics of these variables. The high substitution effects imparted by learning dynamics implies a lower relative standard deviation than in the data.

Table 2: Business Cycle Statistics			
	Data	Model	
		REE	Learning
Panel A: HP Filtered Data			
Relative Std Dev			
σ_A	–	1.21	1.04
σ_C/σ_Y	0.55	0.54 (0.01)	0.43 (0.05)
σ_I/σ_Y	2.88	2.43 (0.02)	2.83 (0.20)
σ_N/σ_Y	0.92	0.49 (0.01)	0.63 (0.07)
σ_{Pr}/σ_Y	0.52	0.54 (0.01)	0.43 (0.05)
Correlation			
$\rho_{C,Y}$	0.78	0.97 (0.01)	0.91 (0.05)
$\rho_{I,Y}$	0.90	0.99 (0.00)	0.98 (0.00)
$\rho_{H,Y}$	0.85	0.97 (0.01)	0.96 (0.01)
$\rho_{Pr,Y}$	0.40	0.97 (0.00)	0.91 (0.05)
$\rho_{Pr,H}$	–0.12	0.88 (0.02)	0.77 (0.08)
Panel B: Growth Rates			
Relative Std Dev			
$\sigma_{\Delta \Delta Y}$	3.96	4.76 (0.26)	4.06 (0.41)
$\sigma_{\Delta C}/\sigma_{\Delta Y}$	0.55	0.51 (0.01)	0.53 (0.08)
$\sigma_{\Delta I}/\sigma_{\Delta Y}$	2.56	2.46 (0.02)	2.60 (0.20)
$\sigma_{\Delta N}/\sigma_{\Delta Y}$	0.80	0.50 (0.00)	0.57 (0.07)
$\sigma_{\Delta Pr}/\sigma_{\Delta Y}$	0.74	0.51 (0.01)	0.53 (0.08)
Correlation			
$\rho_{\Delta C, \Delta Y}$	0.5	0.99 (0.00)	0.90 (0.03)
$\rho_{\Delta I, \Delta Y}$	0.74	0.99 (0.00)	0.96 (0.01)
$\rho_{\Delta N, \Delta Y}$	0.68	0.98 (0.01)	0.91 (0.05)
$\rho_{\Delta Pr, \Delta Y}$	0.62	0.99 (0.01)	0.90 (0.03)
$\rho_{\Delta Pr, \Delta N}$	–0.16	0.94 (0.02)	0.65 (0.07)
Serial Correlation			
ΔC	0.27	0.07 (0.09)	–0.10 (0.10)
ΔI	0.35	–0.03 (0.08)	0.35 (0.07)
ΔY	0.30	–0.00 (0.08)	0.19 (0.08)
ΔH	0.41	–0.03 (0.08)	0.41 (0.07)
ΔPr	–0.06	0.07 (0.09)	–0.10 (0.10)

Note: Pr denotes labor productivity

Panel B shows the same set of statistics in terms of growth rates, underscoring that the model under learning delivers a better fit. In particular, the model does not display the counter factually large output growth volatility which occurs under rational expectations.²⁰ Turning to the correlations between each series and output, all moments are

²⁰The rational expectations model over-predicts the standard deviation of output growth by some 20

closer to the data than are those under rational expectations. Of particular note are the weaker correlations of consumption, wages and labor productivity with output, which are the result of endogenous shifts in the labor supply in response to revisions in beliefs.

Since Cogley and Nason (1993, 1995) and Julio Rotemberg and Michael Woodford (1996), the internal propagation mechanisms of technology shocks have been a central preoccupation of real business cycle theory. The final block of Table 2 reports the autocorrelation properties of the growth rate of key model variables. Investment, output and hours growth are remarkably well matched, even before taking into account the non-negligible sampling error.²¹ As expected, the predictions of the model under rational expectations reveal no internal propagation.

That wages, labor productivity and consumption are counter factually predicted to have negative serial autocorrelation stems from the well-known comovement problem in real business cycle theory emphasized by Robert Barro and Robert G. King (1984) and more recently by Beaudry and Portier (2007). While the impact effect of technology shocks does induce positive comovement, subsequent dynamics under learning are driven by revisions to beliefs which impart positive serial autocorrelation in hours and concomitant negative serial autocorrelation in consumption (statistically different for the reported sampling error). Section 5.4 introduces an extension to the baseline model that resolves these counter factual predictions.

C. Distributions of Beliefs

Because beliefs are central to our story it is useful to study their properties further. Consider the following thought experiment. An econometrician observes an economy with data generated according to the real business cycle model under rational expectations. For each observed sample, the econometrician runs the exact regressions that comprise the beliefs in the learning model — equations (8) – (10) — calibrated with a gain equal to $g = 0.002$. The coefficients are recorded for many simulations.²²

The dashed line in Figure 4 plots a kernel estimate of the implied distribution of the resulting parameter estimates. Six distributions are reported corresponding to the intercept and slope coefficient in each of the three forecasting equations. Because the econometrician is outside the model — equivalently, the econometrician is small relative to the population of rational expectation agents — the distribution reflects pure sampling error: there is no feedback of this sampling error on the true data generating process. The distributions are centered on the rational expectations equilibrium, exhibit negligible bias, and have a fairly small variance. This variance would go to zero as the gain parameter goes to zero, as this would imply that all data are given equal weight. But with the chosen positive gain it is evident that the econometrician has fairly accurate estimates of the parameters characterizing the true data generating process, and would therefore make comparably good forecasts of future returns as the rational agent.

Now imagine a world where all agents modeled by our real business cycle theory actually construct forecasts based on these estimated models. This is precisely the model discussed in this paper. The kernel estimate of the resulting ergodic distribution of the estimated parameters is given by the solid lines. The distribution of the estimated coefficients on capital is not centered on the rational expectations parameters. The distributions are re-centered around the rational expectations coefficients to facilitate comparison with the

percent in contrast to 3 percent for the learning model.

²¹However, higher-order autocorrelations coincide with rational expectations.

²²To compute the distribution of beliefs, the model is simulated 2250 periods and agents' estimates are recorded after discarding the first 2000 observations. The simulation is repeated 250,000 times.

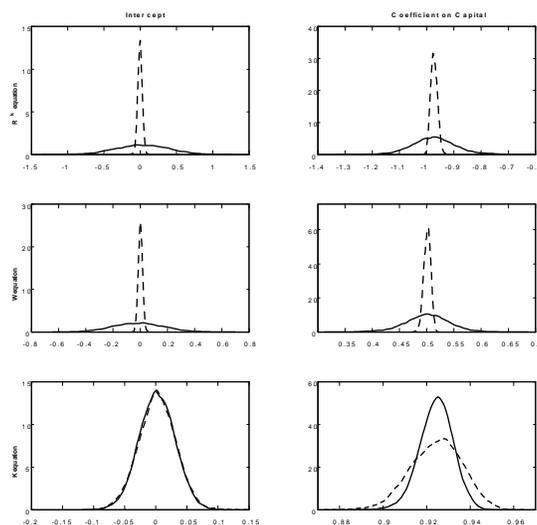


FIGURE 4. DISTRIBUTION OF AGENTS' REGRESSION COEFFICIENTS IN THE MODEL WITH FEEDBACKS (SOLID LINE) AND WITHOUT FEEDBACKS (DOTTED LINE).

non-feedback case.²³ However the median impact impulse responses shown in the previous section indicate that agents' median forecast is in line with rational expectations.

The variation in possible beliefs that can be held by agents is substantially higher than in the previous thought experiment. This dispersion is what leads to the nonlinear impulse response functions and the associated uncertainty of their paths. This in turn generates the increased volatility in the learning model. The figures show that the bulk of the dispersion in agents' beliefs is endogenously determined by the interaction between observed prices and updating of agents' beliefs. This model feature is further manifestation of shifting expectations as a source of business cycle fluctuations that is very much in the spirit of Pigou and Keynes.

D. Generating Comovement

These dynamics relate to a number of recent papers on news shocks and business cycle dynamics — see for example Paul Beaudry and Franck Portier (2007) and Nir Jaimovich and Sergio Rebelo (2008). The present analysis is distinct in the sense that there is only a single source of disturbance — technology shocks. Nonetheless, the negative comovement problem between consumption and hours is manifestation of a common difficulty. Both class of models predict for a given production frontier and separable preferences that non-TFP disturbances engender negative comovement. However, the precise mechanisms are distinct. In the news literature, signals about future productivity generate strong wealth effects that lead to higher consumption and lower labor supply. In

²³The “bias” in the estimates, a product of the nonlinearity of beliefs and linear regression methods, is about 4% for each coefficient. The estimates for the intercept have no bias. As the gain goes to zero this bias vanishes.

the model with learning, the negative correlation between consumption and hours growth stems from strong intertemporal substitution effects. Following Stefano Eusepi and Bruce Preston (2009), this problem is resolved by adopting preferences that are non-separable but consistent with a long-run balanced growth path. We consider an utility function of the form

$$U(C, L) = \frac{(C)^{1-\sigma} v(1-L)}{1-\sigma}, \quad v', v'' > 0, \quad \text{and } \sigma > 1.$$

Nonseparability induces a rise the marginal utility of consumption when labor supply is high, delivering tighter comovement between these variables. As shown in Florin Bilbiie (2009), comovement under this preference specification is obtained only by assuming that consumption is an inferior good. The appendix describes microfoundations with costly labor market participation in which individual household preferences have consumption and leisure being normal goods but in which aggregate dynamics are the same as in this simple representative agent model. Stefano Eusepi and Bruce Preston (2009) develop theoretical implications of a related model in detail.

Following Beaudry and Portier (2007) and Stefano Eusepi (2008), a production technology with variable capacity utilization and a small degree of increasing returns is introduced. The production technology becomes

$$Y_t = \Psi_t (U_t K_t)^\alpha (X_t H_t)^{1-\alpha} \quad \text{where } \Psi_t = \left[(U_t K_t)^\alpha (X_t H_t)^{1-\alpha} \right]^\eta X_t^{-\eta}.$$

The term Ψ_t denotes the external effects of aggregate capital, indexed by the constant $\eta \geq 0$. The term $X_t^{-\eta}$ guarantees that a balanced growth path exists in this model. U_t is the utilization rate of capital in any period t . Capital depreciation is now assumed to increase with capacity utilization according to the function $\delta(U_t) = \theta^{-1} U_t^\theta$. These model features affect firms' labor demand in different ways. Capacity utilization mitigates the decreasing returns to labor, producing a more elastic labor demand, while production externalities induce endogenous labor-demand shifts in response to a changes in expectations. Capacity utilization and increasing returns to scale alone are neither necessary nor sufficient for obtaining comovement in hours and consumption, but they improve the empirical fit of the model and, given the assumed preferences, induce tighter comovement through two channels. First, by amplifying intertemporal substitution of consumption and labor they mute the counterfactual increase in consumption volatility implied by nonseparable preferences. Second, they induce a stronger response of labor demand to shifts in expectations, strengthening comovement.

The assumptions $\sigma = 1$, $\eta = 0$ and $U_t = 1$ for all t delivers our benchmark model. Further details are found in the on-line appendix. There are two extra parameters with respect to the benchmark model. The first parameter, measuring the aggregate externality, is set as $\eta = 0.1$, consistent with the lowest estimate in Marianne Baxter and Robert King (1991). This value implies a "small" degree of externality and a locally determinate equilibrium under rational expectations.²⁴ The second parameter is the household's intertemporal elasticity of substitution, σ , which is chosen to make the ratio of the standard deviations of consumption and output in the model close to the HP-filtered data. This gives $\sigma = 1.5$. The parameter σ_A is again calibrated to match the standard deviation of output in the filtered data. The gain is now $g = 0.001$, half that in the benchmark model. The appendix shows the parameter, θ , indexing variable depreciation, is pinned down by the steady-state return on capital and the steady-state depreciation rate.

²⁴The parameter implies a downward-sloping demand for labor. For the connection between externality and indeterminacy, see Benhabib and Farmer (1994).

Table 3: Business Cycle Statistics				
Data	Nonseparable Preferences		Benchmark	
	REE	Learning	Learning	
Panel A: HP Filtered Data				
Relative Standard Deviation				
σ_A	–	0.71	0.59	1.04
σ_C/σ_Y	0.55	0.66 (0.00)	0.55 (0.03)	0.43 (0.05)
σ_I/σ_Y	2.88	2.15 (0.00)	2.32 (0.09)	2.83 (0.20)
σ_H/σ_Y	0.92	0.56 (0.00)	0.64 (0.04)	0.63 (0.07)
σ_{Pr}/σ_Y	0.52	0.45 (0.00)	0.37 (0.04)	0.43 (0.05)
Correlation				
$\rho_{C,Y}$	0.78	0.99 (0.00)	0.99 (0.00)	0.91 (0.05)
$\rho_{I,Y}$	0.90	0.99 (0.00)	0.99 (0.00)	0.98 (0.00)
$\rho_{H,Y}$	0.85	0.99 (0.00)	0.99 (0.00)	0.96 (0.01)
$\rho_{Pr,Y}$	0.40	0.97 (0.00)	0.97 (0.01)	0.91 (0.05)
$\rho_{Pr,H}$	–0.12	0.95 (0.01)	0.92 (0.02)	0.77 (0.08)
Panel B: Growth Rates				
Relative Standard Deviation				
$\sigma_{\Delta Y}$	3.96	4.84 (0.27)	4.26 (0.54)	4.06 (0.41)
$\sigma_{\Delta C}/\sigma_{\Delta Y}$	0.55	0.61 (0.00)	0.59 (0.04)	0.53 (0.08)
$\sigma_{\Delta I}/\sigma_{\Delta Y}$	2.56	2.16 (0.00)	2.23 (0.11)	2.60 (0.20)
$\sigma_{\Delta N}/\sigma_{\Delta Y}$	0.80	0.56 (0.00)	0.6 (0.05)	0.57 (0.07)
$\sigma_{\Delta Pr}/\sigma_{\Delta Y}$	0.74	0.45 (0.00)	0.43 (0.06)	0.53 (0.08)
Correlation				
$\rho_{\Delta C,\Delta Y}$	0.5	0.99 (0.00)	0.99 (0.00)	0.90 (0.03)
$\rho_{\Delta I,\Delta Y}$	0.74	0.99 (0.00)	0.99 (0.00)	0.96 (0.01)
$\rho_{\Delta N,\Delta Y}$	0.68	0.99 (0.00)	0.98 (0.00)	0.91 (0.05)
$\rho_{\Delta Pr,\Delta Y}$	0.62	0.96 (0.00)	0.96 (0.01)	0.90 (0.03)
$\rho_{\Delta Pr,\Delta N}$	–0.16	0.98 (0.01)	0.89 (0.03)	0.65 (0.07)
Serial Correlation				
ΔC	0.27	–0.01 (0.08)	0.08 (0.08)	–0.10 (0.10)
ΔI	0.35	–0.03 (0.08)	0.25 (0.08)	0.35 (0.07)
ΔY	0.30	–0.03 (0.08)	0.17 (0.08)	0.19 (0.08)
ΔH	0.41	–0.03 (0.08)	0.30 (0.08)	0.41 (0.07)
ΔPr	–0.06	0.01 (0.08)	–0.01 (0.08)	–0.10 (0.10)

Note: Pr denotes labor productivity

Table 3 reports an analogous set of statistics to table 2 for the separable preferences model. Space constraints prohibit an exhaustive discussion, but the following are pertinent. The model does well in most dimensions. It generates a 16% increase in output volatility compared to rational expectations. Non-separable preferences achieve a stronger correlation between consumption and hours, reflected in the positive autocorrelation of the former. This comes at the cost of slightly lower volatility of investment relative to the benchmark model. These results address some of the concerns regarding predictable movements laid out in Rotemberg and Woodford (1996). Finally, the model continues to

generate forecast errors consistent with the documented patterns implied by data from the Survey of Professional Forecasters.

V. Robustness

This section provides two additional robustness exercises. Unless otherwise noted, parameters are held fixed at benchmark values for the model under learning. The standard deviation of technology shocks remains unchanged across simulations.

Alternative gain parameters. Table 4 reports statistics on amplification, propagation and serial correlation in model-implied forecast errors as the gain parameter is varied. In addition to our benchmark model, which is reproduced in column 4, six new gains are considered spanning the interval (0.0013, 0.003). Amplification is judged by the difference between the output standard deviations in the learning and rational expectations models. We also report the volatility of investment and hours (relative to output) that is due to learning. Propagation is judged by the serial correlation in output growth in the learning model. The final row, Forecast errors, gives the serial correlation in the model one-period-ahead forecast errors for both interest rates and wages. A single number is reported as the statistics are identical under maintained parametric assumptions. For this range of gain parameters the model continues to produce amplification and propagation of technology shocks, while also being consistent with the documented patterns of forecast errors in survey expectations data. Smaller gains render predictions closer to rational expectations; but even the smallest gain provides some amplification of technology shocks, 10 percent, and persistence in forecast errors of 0.22. Larger gains provide non-trivial amplification, propagation and serial correlation. The largest gain, which remains conservative, produces amplification in output, investment and hours volatility of 20, 25 and 43 percent respectively and persistence in forecast errors of 0.52 — slightly below the largest statistic in the Survey of Professional Forecast data for unemployment rates of 0.58.

	Gain						
	0.0013	0.0015	0.0017	0.002	0.0025	0.0027	0.003
Amplification:							
σ_Y	9.7%	11%	12%	14%	17%	18%	20%
σ_I/σ_Y	11%	12%	14%	17%	21%	23%	25%
σ_H/σ_Y	18%	21%	24%	28%	36%	39%	43%
Propagation	0.12	0.14	0.16	0.19	0.23	0.25	0.27
Forecast errors	0.22	0.26	0.3	0.36	0.45	0.48	0.52

Note: the table shows the properties of the benchmark model under different values of the constant gain g . The first three rows show the excess volatility in HP detrended output, investment and hours with respect rational expectations. The fourth row shows the first order autocorrelation of output growth. The final row displays the autocorrelation of one-quarter-ahead forecast errors. The reported statistics from the model are mean values over 5000 simulations of length 162 quarters (excluding 2000 quarters of initial simulated data that are discarded). The numbers in parenthesis denote the standard deviation across simulations.

Alternative models. This section considers alternative model assumptions. Results are collated in table 5. Models 1 and 2 show the benchmark results for the rational

expectations and learning models. The latter reiterates earlier results for ease of comparison while the former gives the results under rational expectations assuming the same standard deviation of technology shocks as model 2. The improved amplification is again immediate. Models 3 and 4 show the cases of a lower elasticity of labor supply ($\epsilon_H = 0.5$). Under both rational expectations and learning, the volatility of output falls for a given standard deviation technology shock. Concomitantly, the relative volatility of investment and hours also decline, while the relative volatility of consumption increases. The serial correlation properties adjust accordingly. These results underscore the centrality of the elasticity of labor supply in generating plausible volatility in real business cycle models.

Table 5: Alternative model specifications.

	Statistics							
	σ_Y	σ_C/σ_Y	σ_I/σ_Y	σ_H/σ_Y	Δ_C	Δ_Y	Δ_I	$(FE_t^{1Q}, FE_{t-1}^{1Q})$
Data	1.54	0.52	2.82	0.91	0.16	0.34	0.27	–
1. Baseline RE	1.31	0.54	2.43	0.49	0.07	–0.00	–0.03	–0.01 (0.08)
2. Baseline Learn	1.54	0.43	2.83	0.63	–0.1	0.19	0.35	0.36 (0.19)
3. Low Elast. RE	0.9	0.58	2.29	0.3	0.07	0.00	–0.02	–0.01 (0.08)
4. Low Elast. Learn	1.13	0.51	2.56	0.36	–0.13	0.11	0.30	0.09 (0.14)
5. Inv. Adj. costs: $\phi'' = 0.5$	1.06	0.68	2.09	0.39	–0.04	0.16	0.56	–0.01 (0.08)
6. Euler Equation: $\gamma = 0.002$	1.31	0.54	2.43	0.49	0.08	–0.01	–0.03	0.00 (0.08)
$\gamma = 0.04$	1.31	0.54	2.43	0.49	0.08	0.00	–0.02	0.09 (0.07)

Model 5 presents a model under rational expectations with investment adjustment costs. This permits a comparison of learning dynamics with one popular friction employed in the real business cycle literature. A more exhaustive comparative exercise is beyond the scope of this paper. Introducing investment adjustment costs of the form

$$K_{t+1} = I_t \left[1 - \phi \left(\frac{I_t}{I_{t-1}} \right) \right] + (1 - \delta)K_t$$

with $\phi(\bar{\gamma}) = \phi'(\bar{\gamma}) = 0$ and $\phi''(\bar{\gamma}) > 0$ in the rational expectations model certainly improves correspondence of model predictions with data on some dimensions — the first-order serial correlation properties of output and investment are much improved and output is more volatile.²⁵ But remaining moments are, if anything, further from the data. In particular, the relative volatility of investment is considerably dampened.

The final rows report statistics for an alternative model with learning. Many recent papers have proposed analyses of learning dynamics in the context of models where agents solve finite-horizon decision problems, but without requiring that agents make forecasts more than one period into the future. Agents' decisions depend only on forecasts of future variables that appear in Euler equations used to characterize rational expectations

²⁵In this experiment $\phi''(\bar{\gamma})$ are chosen to obtain the closest first order autocorrelation of output growth to the data.

equilibrium. Key contributions include James Bullard and Kaushik Mitra (2002) and George W. Evans and Seppo Honkapohja (2003).

Of particular relevance to the present study are the analyses of Noah Williams (2003) and Eva Carceles-Poveda and Chryssi Giannitsarou (2007). The former studies precisely the question explored here: can learning be a source of business cycle fluctuations? The latter is similarly motivated, with specific focus on asset pricing implications of real business cycle theory. Both papers make use of models with learning dynamics in which only one-period-ahead expectations matter to expenditure and production plans of households and firms. Both conclude that learning of the kind considered here is unpromising in generating amplification and propagation.²⁶

The final two rows replicate this kind of analysis in the context of the model developed here. The two models are differentiated by choice of gain parameter. Williams (2003) proceeds assuming that the Euler equations predicted by a rational expectations analysis of the model represent decision rules of agents under learning. The model under learning then assumes household consumption decisions are determined as

$$(17) \quad c_t = \hat{E}_t c_{t+1} - \hat{E}_t (\beta \bar{R} R_{t+1}^K + \hat{\gamma}_{t+1}).$$

This requires the further assumption that households directly forecast their own future consumption using regressions of the kind specified in section 2. Preston (2005) shows that this decision rule leads to suboptimal decisions — see also Marcet and Sargent (1989).²⁷ All remaining model equations are unchanged as they do not directly depend on the specification of beliefs.

This approach leads to dramatically different conclusions. Learning dynamics fail to generate amplification and propagation, for either a model with our benchmark gain ($\gamma = 0.002$) or a model with significantly larger gain ($\gamma = 0.04$). Model-implied moments are essentially indistinguishable from a rational expectations analysis of the model; though the large-gain model generates some persistence in expectational errors. That learning models with only one-period-ahead expectations fail to capture salient features of data relative to models based on optimal decision rules is also supported by Sophocles Mavroeidis and Guillaume Chevillon and Michael Massmann (2009).

This negative finding has less to do with learning than it does with the assumed nature of economic decisions. In real business cycle theory the only intertemporal decision is the household's consumption and saving decision. To make this decision households must forecast the entire future sequence of wages and real interest rates. These beliefs about future prices determine current market clearing prices, which in turn determine beliefs. A consequence of the model of household behavior given by (17) is the connection between market prices that govern future consumption and investment opportunities and current allocations and prices is severed. The economic structure of the model is completely changed and revealed to matter greatly for implied model dynamics. The strength of the approach followed in this paper is that agents' decision rules are consistent with the model's microfoundations.²⁸

²⁶A final related paper is Kevin Huang, Zheng Liu and Tao Zha (2008). It considers the same model as Williams (2003) where only one-period-ahead expectations matter, but examines a belief structure that does not nest the rational expectations equilibrium of the model. However, no attempt is made to calibrate the model to fit observed data and results depend on one specific choice of initial beliefs.

²⁷That (17) describes optimal decisions under rational expectations and not learning reflects the property under rational expectations of equilibrium probability laws embedding information about all relevant constraints, including transversality conditions and intertemporal budget constraints. This is not true once beliefs are exogenously specified as in the learning model contemplated here.

²⁸The two models also differ in their local stability properties under learning dynamics. Consistently with James Bullard and Stefano Eusepi (2008), models with infinite-horizon decisions imply larger eigen-

More specifically, the on-line appendix demonstrates that in the period after a technology shock, revised beliefs are identical under the optimal decision and Euler equation models — recall that in either case beliefs are centered on the rational expectations equilibrium on average, and that under our benchmark calibration consumption and wages are essentially identical; the forecasting models in each case are therefore also identical. Moreover, revised beliefs imply forecasts at short horizons that are almost identical to rational expectations forecasts. However, they are different at long horizons. This reflects the fact that variation in capital dominates forecasts at short horizons but not long horizons due to reversion to steady state. At long horizons it is variations in belief parameters that cause departures from rational expectation forecasts. Because the Euler equation approach does not depend on these forecasts it fails to give predictions that differ much from rational expectations equilibrium.

VI. Conclusion

In the spirit of Pigou (1927) a model with learning dynamics is developed in which self-fulfilling expectations are possible in response to technology shocks. The benchmark model delivers volatility in output comparable to a rational expectations analysis with a standard deviation of technology shock that is 10 - 20 percent smaller, and has substantially more volatility in investment and hours. The model captures persistence in these series, unlike standard models. The improvement in fit stems from shifting beliefs amplifying standard income and substitution effects operative in real business cycle theory. An important feature of the model is its ability to replicate patterns in forecast errors over the business cycle implied by data from the Survey of Professional Forecasters.

Data Appendix

We use data for the US economy, 1948:Q1 to 2007:Q4. The variables are constructed as follows (DLX codes in parenthesis). Output is Real Gross Domestic Product (GDPH); nominal consumption is computed as the sum of nondurable goods (CN), services (CS) and government expenditures (G); nominal investment is the sum of private nonresidential investment structures (FNS), Equipment and software (FNE), private residential investment (FR) and consumption durable goods (CD). Consumption and investment are converted in real terms by using the GDP deflator (GDP/GDPH). For total hours we use the measure by Francis and Ramey (2008). All variables are transformed in per capita terms by using the civilian non-institutional population between 22 and 64 years old, also from Francis and Ramey (2008). Productivity is measured as real GDP divided by total hours worked.

Survey forecast data is from the Survey of Professional Forecasters, collected by the Federal Reserve Bank of Philadelphia. The median one- and four-quarter-ahead forecasts are used for a range of macroeconomic variables. The survey is quarterly; it is sent out to participants at the end of the first month of each quarter and response deadlines are the middle month of each quarter. A detailed description of the data set can be found at the web site: <http://www.phil.frb.org/research-and-data/real-time-center/survey-of-professional-forecasters/>. The available forecasts cover sample periods of variable lengths. Forecasts for Nominal GDP, Industrial Production, the unemployment rate and GDP deflator are available for 1968Q4-2009Q1. Forecasts for Real GDP, the T-Bill (3 month), Corporate Bond and CPI are available for 1981Q3-2009Q1. The Tbond (10 years) is available for 1992Q1-2009Q1. We consider two additional series. The

values than models under one-period-ahead decisions, affecting the convergence properties of the model under learning.

first is a longer series for real GDP which is obtained by combining forecasts of nominal GDP with forecast of the GDP deflator (both available starting in 1968Q4). The second series is a forecast of the expected (ex-post) real interest rate, obtained by subtracting the two-period-ahead forecast of GDP deflator from the one-period ahead forecast of the Tbill nominal interest rate. With the exception of the (ex-post) real interest rate, we consider one- and four-period-ahead forecasts. We compute forecast errors as the difference between forecasts and the realized variable. The realized variable corresponds to the latest available vintage of the data. For robustness, for real GDP (extended sample) we compute the one-quarter-ahead forecast errors using alternative data vintages available to forecasters. In particular we used the first, second, fifth and ninth vintage available.

REFERENCES

- Andolfatto, David.** 1996. "Business Cycles and Labor-Market Search." *American Economic Review*, 86(1): 112.
- Barro, Robert, and Robert G. King.** 1984. "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles." *Quarterly Journal of Economics*, 817–839.
- Baxter, Marianne, and Robert King.** 1991. "Productive Externalities and Business Cycles." manuscript.
- Baudry, Paul, and Franck Portier.** 2007. "When Can Changes in Expectations Cause Business Cycle Fluctuations in Neo-Classical Settings?" *Journal of Economic Theory*, 135: 458–477.
- Benhabib, Jess, and Roger Farmer.** 1994. "Indeterminacy and Increasing Returns." *Journal of Economic Theory*, 63: 19–41.
- Bilbiie, Florin.** 2009. "Nonseparable preferences, fiscal policy puzzles and inferior goods." *Journal of Money, Credit and Banking*.
- Branch, William A., and George W. Evans.** 2006. "A Simple Recursive Forecasting Model." *Economics Letters*, (91): 158–166.
- Bullard, James, and Kaushik Mitra.** 2002. "Learning About Monetary Policy Rules." *Journal of Monetary Economics*, 49(6): 1105–1129.
- Bullard, James, and Stefano Eusepi.** 2008. "When Does Determinacy Imply Learnability?" Federal Reserve Bank of St. Louis Working Paper (007A).
- Burnside, Craig, and Martin Eichenbaum.** 1996. "Factor-Hoarding and the Propagation of Business-Cycle Shocks." *American Economic Review*, 86(5): 1154–1174.
- Carceles-Poveda, Eva, and Chryssi Giannitsarou.** 2007. "Asset Pricing with Adaptive Learning." forthcoming, *Review of Economic Dynamics*.
- Carlstrom, Charles T., and Timothy S. Fuerst.** 1997. "Agency Cost, Net Worth, and Business Cycle Fluctuations: A Computable General Equilibrium Analysis." *American Economic Review*, 87(5): 893–910.
- Christiano, Lawrence J., and Martin Eichenbaum.** 1992. "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations." *American Economic Review*, 82: 430–450.
- Cogley, Timothy, and James M. Nason.** 1993. "Impulse Dynamics and Propagation Mechanisms in a Real Business Cycle Model." *Economics Letters*, 453: 77–81.
- Cogley, Timothy, and James M. Nason.** 1995. "Output Dynamics in Real-Business-Cycle Models." *American Economic Review*, 85(3): 492–511.
- Edge, Rochelle, Thomas Laubach, and John Williams.** 2007. "Learning and Shifts in Long-Run Productivity Growth." *Journal of Monetary Economics*, 2421–2438.

- Eusepi, Stefano.** 2008. "A Note on Expectations-Driven Business Cycles in Economies with Production Externalities." forthcoming *International Journal of Economic Theory*.
- Eusepi, Stefano, and Bruce Preston.** 2009. "Labor Supply Heterogeneity and Macroeconomic Co-movement." NBER Working Paper no. 15561.
- Evans, George W., and Seppo Honkapohja.** 2001. *Learning and Expectations in Economics*. Princeton, Princeton University Press.
- Evans, George W., and Seppo Honkapohja.** 2003. "Expectations and the Stability Problem for Optimal Monetary Policies." *Review of Economic Studies*, 70(4): 807–824.
- Hall, Robert E.** 1988. "Intertemporal Substitution in Consumption." *Journal of Political Economy*, 96.
- Hansen, Gary.** 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics*, 16: 309–337.
- Huang, Kevin, Zheng Liu, and Tao Zha.** 2008. "Learning, Adaptive Expectations, and Technology Shocks." unpublished, Federal Reserve Bank of Atlanta.
- Jaimovich, Nir, and Sergio Rebelo.** 2008. "Can News About the Future Drive the Business Cycle." unpublished, Northwestern University.
- King, Robert G., Charles I. Plosser, and Sergio T. Rebelo.** 1988. "Production, Growth and Business Cycles." *Journal of Monetary Economics*, 21: 195–232.
- Kreps, David.** 1998. "Anticipated Utility and Dynamic Choice." In *Frontiers of Research in Economic Theory*, ed. D.P. Jacobs, E. Kalai and M. Kamien, 242–274. Cambridge: Cambridge University Press.
- Kydland, Finn E., and Edward C. Prescott.** 1982. "Time to Build and Aggregate Fluctuations." *Econometrica*, 50(6): 1345–1370.
- Lorenzoni, Guido.** 2008. "A Theory of Demand Shocks." unpublished, MIT.
- Marcet, Albert, and Thomas J. Sargent.** 1989. "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information." *Journal of Political Economy*, 1306–1322.
- Mavroeidis, Sophocles, Guillaume Chevillon, and Michael Massmann.** 2010. "Inference in models with adaptive learning." forthcoming, *Journal of Monetary Economics*.
- Milani, Fabio.** 2007. "Expectations, Learning and Macroeconomic Persistence." *Journal of Monetary Economics*, 54: 2065–2082.
- Orphanides, Athanasios, and John C. Williams.** 2005. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." University of Chicago Press.
- Piazzesi, Monika, and Martin Schneider.** 2008. "Bond Positions, Expectations, and the Yield Curve." unpublished, Stanford University.
- Pigou, Alfred C.** 1927. *Industrial Fluctuations*. London: MacMillan.
- Prescott, Edward.** 1986. "Theory Ahead of Business Cycles Measurement." *Carnegie-Rochester Conference Series on Public Policy*, 25: 11–66.
- Preston, Bruce.** 2005. "Learning About Monetary Policy Rules when Long-Horizon Expectations Matter." *International Journal of Central Banking*, 1(2): 81–126.
- Roberts, John M.** 1997. "Is Inflation Sticky?" *Journal of Monetary Economics*, 39: 173–196.
- Rogerson, Richard.** 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics*, 21: 3–16.
- Rotemberg, Julio, and Michael Woodford.** 1996. "Real Business Cycle Models and the Forecastable Movements in Output, Hours and Consumption." *American Economic*

Review, 86(1): 71–89.

Sargent, Thomas J. 1999. *The Conquest of American Inflation*. Princeton University Press.

Schmitt-Grohe, Stephanie. 2000. “Endogenous Business Cycles and the Dynamics of Output, Hours and Consumption.” *American Economic Review*, 90: 1136–1159.

Thomas, Lloyd B. 1999. “Survey Measures of Expected U.S. Inflation.” *Journal of Economic Perspectives*, 125–144.

Williams, Noah. 2003. “Adaptive Learning and Business Cycles.” unpublished, Princeton University.