

Hommes, Cars H.; Mavromatis, Kostas; Özden, Tolga; Zhu, Mei

Article

Behavioral Learning Equilibria in New Keynesian models

Quantitative Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Hommes, Cars H.; Mavromatis, Kostas; Özden, Tolga; Zhu, Mei (2023) : Behavioral Learning Equilibria in New Keynesian models, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 14, Iss. 4, pp. 1401-1445, <https://doi.org/10.3982/QE1533>

This Version is available at:

<https://hdl.handle.net/10419/296347>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

Behavioral learning equilibria in New Keynesian models

CARS HOMMES

Bank of Canada and CeNDEE, Amsterdam School of Economics, University of Amsterdam

KOSTAS MAVROMATIS

De Nederlandsche Bank and Amsterdam School of Economics, University of Amsterdam

TOLGA ÖZDEN

Bank of Canada

MEI ZHU

Institute for Advanced Research and School of Economics, Shanghai University of Finance and Economics
and The Key Laboratory of Mathematical Economics (SUFE), Ministry of Education

We introduce Behavioral Learning Equilibria (BLE) into a multivariate linear framework and apply it to New Keynesian DSGE models. In a BLE, boundedly rational agents use simple, but optimal AR(1) forecasting rules whose parameters are consistent with the observed sample mean and autocorrelation of past data. We study the BLE concept in a standard 3-equation New Keynesian model and develop an estimation methodology for the canonical Smets and Wouters (2007) model. A horse race between Rational Expectations (REE), BLE, and constant gain learning models shows that the BLE model outperforms the REE benchmark and is competitive with constant gain learning models in terms of in-sample and out-of-sample fitness. Sample-autocorrelation learning of optimal AR(1) beliefs provides the best fit when short-term survey data on inflation expectations are taken into account in the estimation. As a policy application, we show that optimal Taylor rules under AR(1) expectations inherit history dependence and require a lower degrees of interest rate smoothing than REE.

KEYWORDS. Bounded rationality, adaptive learning, estimation, behavioral New Keynesian macro model, monetary policy.

JEL CLASSIFICATION. C11, D83, D84, E03, E62.

Cars Hommes: C.H.Hommes@uva.nl

Kostas Mavromatis: K.Mavromatis@dnb.nl

Tolga Özden: tozden@bank-banque-canada.ca

Mei Zhu: Zhu.Mei@mail.shufe.edu.cn

The views expressed in this paper do not represent the positions of the Bank of Canada, De Nederlandsche Bank, or the Eurosystem. Stimulating discussions and comments from Klaus Adam, Bill Branch, Jim Bullard, George Evans, Stephanie Schmitt-Grohé, Paul Levine, Domenico Massaro, Bruce McGough, Bruce Preston, Sergey Slobodyan, Gianluca Violante, Mike Woodford, Raf Wouters, and the anonymous referees are gratefully acknowledged. Mei Zhu acknowledges financial support from NSFC (11401365, 72141303).

1. INTRODUCTION

Rational expectations (RE) is the workhorse approach for modeling expectations in DSGE models, and it has been the dominant framework in macroeconomic modeling for several decades since the work of [Muth \(1961\)](#) and [Lucas \(1972\)](#). The RE paradigm is a model-consistent approach where, by construction, agents' expectations are on average confirmed by the realizations of the model. Nevertheless, some drawbacks of RE models have been highlighted in recent literature. One of these shortcomings is matching the persistence of macroeconomic variables. To do so, RE models typically need to be augmented by highly persistent exogenous shocks or other sources of persistence such as consumption habits and indexation in prices and wages. Agents in RE models are assumed to know a large number of state variables, shocks, and parameters to form their expectations. In medium- and large-scale DSGE models, such assumptions lead to implausibly large information sets. Some studies have also highlighted the failure of RE models to match expectations data from standard surveys ([Coibion, Gorodnichenko, and Kamdar \(2018\)](#)).

In this paper, we propose a Behavioral Learning Equilibrium (BLE) as a plausible and parsimonious alternative to RE that matches persistence and fits with survey data. A BLE is one of the most parsimonious misspecification equilibria, where agents use a simple forecasting model because the economy is too complex to fully understand its structure. Along a BLE, agents forecast the states of the economy by simple, but optimal univariate AR(1) rules.¹ The AR(1) rules are optimal in the sense that the mean and the first-order autocorrelation of all forecasts coincide with the actual mean and the first-order autocorrelation of the realizations.

[Hommes and Zhu \(2014\)](#) applied the BLE concept in the simplest framework of a linear univariate model driven by autocorrelated shocks. In this paper, we extend it to multivariate linear systems and provide a method for approximating and estimating a BLE in a general setup. We use Bayesian methods to estimate BLE in the medium-scale Smets–Wouters (2007) DSGE model and compare the in-sample fit and the out-of sample forecasting performance to the Rational Expectations Equilibrium (REE) benchmark and alternative learning models. An advantage of the BLE model, relative to adaptive learning models, is that it places significantly more restrictions on the agents' forecasting model. As argued in [Gaus and Gibbs \(2018\)](#), adaptive learning models typically achieve an improvement in model fitness by breaking the cross-equation restrictions of the underlying REE model. A BLE disciplines the degree of breaking cross-equation restrictions, but still achieves significant improvement in model fitness.

¹Different types of misspecification equilibria have been proposed in the literature. A nonexhaustive list includes Restricted Perceptions Equilibria (RPE), which generally refer to underparameterized forecasting rules (see, e.g., [Sargent \(1991\)](#), [Evans and Honkapohja \(2001\)](#), [Branch \(2004\)](#), [Adam \(2007\)](#), [Bullard, Evans, and Honkapohja \(2008\)](#), [Lansing \(2009\)](#), [Branch and Evans \(2010\)](#), [Lansing and Ma \(2017\)](#), [Audzei and Slobodyan \(2022\)](#), and Natural Expectations ([Fuster, Laibson, and Mendel \(2010\)](#))) where agents use autoregressive models with lower orders than implied by the correct model. The closest misspecification equilibrium to our work is that of Consistent Expectations Equilibria (CEE) ([Hommes and Sorger \(1998\)](#)), where agents use a simple linear AR(1) rule in a nonlinear model.

One of the appealing features of RE models is that they remove all parameters and degrees of freedom associated with expectations. RE are model-consistent and are determined by the structural parameters. A BLE is also subject to a set of restrictions and, therefore, it is parameter-free and completely pinned down by structural parameters. In this sense, a BLE is an *equilibrium* model where the parameters of the AR(1) rules have been set optimally akin to a REE. The models differ in terms of the information set of agents' knowledge about the underlying system.² In the linearized DSGE framework, REE and BLE are both linear equilibrium models but they satisfy different fixed-point conditions. While REE assumes perfect knowledge of the underlying multivariate linear structure, BLE imposes observable consistency restrictions that the first two moments, the mean and the first-order autocorrelation, must satisfy. These conditions imply that the optimal AR(1) rules are unbiased and their forecast errors are uncorrelated with predictor variables, but these observable restrictions are less strong than the perfect fixed-point conditions for model-consistent REE.

Our paper makes theoretical and empirical as well as policy contributions. In terms of theoretical contributions, we derive existence conditions of BLE in a general linear framework and stability conditions for a natural learning process of BLE, the sample autocorrelation (SAC-)learning. We then apply these results to the simplest New Keynesian (NK) model (Woodford (2003a)), show that the Taylor principle is sufficient for the existence of a BLE and study its E-stability under SAC-learning.

In terms of an empirical application, we use the Smets–Wouters (2007) DSGE model as a test ground for a horse race between BLE, REE, several constant-gain recursive least squares models (pseudo MSV, AR(2), and VAR(1)), and SAC-learning by comparing the models across a multitude of dimensions. In particular, we compare the models in terms of in-sample fitness and pseudo out-of-sample forecasting performance. We further discuss their performance to match short-term inflation expectations by estimating the models with data from the Survey of Professional Forecasters (SPF). We find that the BLE model generally improves upon the REE benchmark in terms of both in-sample fitness and pseudo out-of-sample forecasting performance, while learning models tend to outperform the equilibrium models BLE and REE. Among the learning models, we find that SAC-learning yields the best model fitness and matches short-term inflation survey expectations data well.

In terms of policy application, we investigate optimal smoothing within the class of standard Taylor rules and find that optimal interest rate smoothing is substantially lower in the BLE model than in the REE model. This result extends to SAC-learning, while the pseudo MSV-learning model yields an optimal smoothing degree closer to the REE benchmark. This suggests that when expectations are persistent and backward-looking, as in the case of BLE, the central bank does not need to introduce more persistence and history-dependence through interest rate smoothing, as in the case of REE. We show that the deployment by agents of simple backward-looking rules to forecast macroeconomic aggregates makes the interest rate dependent on past data, and thus adds history

²We introduce BLE by taking the microfoundations of DSGE models as given. We elaborate further on this point in Section 2 where BLE is formally introduced.

dependence in policy rate setting. When agents are purely forward-looking instead, as in REE, interest rate smoothing is necessary in order for policy-rate decisions to become history dependent.³

At the time of the writing of this paper, major central banks like the Federal Reserve and the European Central Bank are reviewing their strategies. The fear of failing to anchor inflation expectations well has led central banks to broaden the range of models used for the analysis of monetary policy transmission. In particular, the analysis of monetary policy transmission is deemed necessary in models where expectations are no longer rational but feature bounded rationality and backward-looking behavior.⁴ This reveals that our analysis lies at the heart of current policy debates since we estimate one of the most prominent models in central banking by accounting for various types of learning as a deviation from the rational expectations benchmark.

The paper is organized as follows. Section 2 focuses on theory. It introduces the main concepts of BLE in a general n -dimensional setup, presents the existence and stability conditions of BLE in a multivariate linear framework, applies BLE in the baseline 3-equation NK model, and presents a numerical method to approximate an E-stable BLE. Section 3 is an empirical application using the Smets–Wouters NK model to run a horse race between different equilibrium and learning models using a Bayesian estimation methodology. Section 4 discusses a policy application of optimal interest rate smoothing, comparing the equilibrium and some of the learning models. Finally, Section 5 concludes.

Related literature

Applications of adaptive learning in macroeconomic models have been of great interest to policymakers and academics alike. Our paper contributes to this growing line of literature; see, for example, [Evans and Honkapohja \(2001\)](#), [Branch and Evans \(2006\)](#), [Bullard \(2006\)](#), [Woodford \(2013\)](#), and [Angeletos and Lian \(2016\)](#) for extensive reviews.⁵

A shortcoming of REE models that has received attention in the literature is their failure to generate realistic expectation dynamics and being at odds with data coming from survey expectations. For example, [Canova and Gambetti \(2010\)](#) revisit the great moderation period and examine the role of expectations using reduced form methods. By using data from SPE, they find an important role for expectations that did not substantially change over time. [Adam and Padula \(2011\)](#) estimate a forward-looking New

³In the literature on the design of optimal monetary policy under rational expectations, history dependence is also obtained through price level targeting instead of inflation targeting. For a more detailed discussion, see [Giannoni \(2014\)](#) and the references therein.

⁴In her speech on September 30, 2020, at the ECB and its watchers XXI conference, Christine Lagarde alluded to the relevance of models that depart from the rational expectations assumption by stating, “*while make-up strategies may be less successful when people are not perfectly rational in their decisions—which is probably a good approximation of the reality we face—the usefulness of such an approach could be examined.*”

⁵There is a large body of literature on the analysis of learning in macroeconomic models (see [Huang, Liu, and Zha \(2009\)](#), [Marcet and Nicolini \(2003\)](#), [Sargent, Williams, and Zha \(2009\)](#) and [Williams \(2003\)](#), among others.) In this paper, we restrict ourselves to the literature on the analysis of monetary policy under learning.

Keynesian Phillips Curve (NKPC) using data from the SPF (Croushore (1993)) as a proxy for expected inflation and obtain reasonable estimates for the slope of the NKPC, which is an improvement over the REE model. Along similar lines, Del Negro and Eusepi (2011) use inflation expectations as an observed variable in their model estimations and find evidence that the survey of expectations contains information not explained by other macroeconomic variables. Gennaioli, Ma, and Shleifer (2016) show, by using survey expectations, that corporate investment plans depend on CFOs' expectations of earnings growth. Forecast errors in CFOs' expectations are predictable, which provides evidence in support of small extrapolative forecasting rules. Fuhrer (2017) shows that embedding survey data into DSGE models helps in several directions, such as reducing reliance on ad hoc sources of persistence like habit and indexation. A common feature in these studies is that they document the shortcomings of REE models along the expectations dimension and argue for the usefulness of incorporating data from survey expectations into these models.

Much of the literature on adaptive learning focuses on dynamics under MSV-learning of a correctly specified model (see, e.g., Marcet and Sargent (1989), Evans and Honkapohja (2001), Milani (2007)) and studies conditions under which the learning process converges on the underlying REE. Orphanides and Williams (2004) study monetary policy under MSV-learning and find that optimal policy is typically more aggressive to inflation under learning. Milani (2007, 2011) considers the estimation of the baseline NK model and finds that the model fit is improved under learning, while the dependence on some structural parameters such as habit and indexation is substantially reduced. Berardi and Galimberti (2017) consider model specifications with time-varying gains under MSV-learning and find higher estimates for the gain parameter on inflation.

In a related study, Gaus and Gibbs (2018) consider models with Euler-equation learning (Evans and Honkapohja (2003)) and infinite-horizon learning (Preston (2005)) to compare with the REE benchmark. They document that introducing adaptive learning in DSGE models leads to a near-universal improvement in model fit, while the estimated parameter bands remain mostly unchanged compared to REE. Gaus and Gibbs (2018) then compare their learning models to *fixed beliefs* (FB) models and show that much of the improved model fit is due to relaxing the cross-equation restrictions of REE. Our approach complements and extends their analysis in several dimensions. First, Gaus and Gibbs (2018) do not consider misspecified rules but use FB with a correctly specified forecasting function (the MSV solution) with fixed parameters, which they set equal to the estimated REE belief parameters. Our BLE concept with an AR(1) forecasting rule is one of the most parsimonious misspecified rules (using only a constant [the mean] and the lagged state variable, and no exogenous shocks). Second, we introduce a *fixed-beliefs equilibrium*, where the parameters of the AR(1) rule are *optimized* using the behavioral restrictions imposed by BLE, namely that the mean and first-order autocorrelations are correct. Hence, we study whether the behavioral equilibrium cross-equation restrictions of a BLE improve the model fit. Third, BLE comes with a natural learning scheme: SAC-learning. Therefore, we can disentangle the empirical fit of the behavioral BLE restrictions and its SAC-learning process and study whether learning adds to improving the empirical fit.

A growing number of papers also consider small and/or misspecified forecasting rules as a convenient alternative to RE and MSV-learning. [Lansing \(2009\)](#) constructs a consistent expectations equilibrium (CEE), similar in spirit to our BLE concept, where agents use the optimal Kalman gain within their class of misspecified models. Along similar lines, [Lansing and Ma \(2017\)](#) use a CEE concept to study exchange rate dynamics. [Fuster, Hebert, and Laibson \(2010, 2012\)](#) and [Fuster, Laibson, and Mendel \(2010\)](#) study natural expectations characterized by an underestimation of the degree of mean reversion, which arises when agents use lower-order autoregressive models than is warranted by the correct data generating process. As such, when applied to models of higher-order autoregressive processes, a BLE may be seen as the simplest case of natural expectations. [Ormeño and Molnár \(2015\)](#) investigate whether an adaptive learning model can fit the macroeconomic and survey data simultaneously and find that this is true only when small forecasting rules are considered. The most relevant study for this paper is [Slobodyan and Wouters \(2012a\)](#), where the authors show that an AR(2) forecasting rule under Kalman gain learning substantially improves the model fit without a large effect on parameter estimates. As such, this paper can be seen as extending their work in several directions, where we disentangle the effects of the fixed-equilibrium beliefs, the timing of expectations and the learning algorithm on the model fit. [Audzei and Slobodyan \(2022\)](#) consider a model where agents use misspecified models, and they are allowed to evaluate and change their forecasting models over time. They find that in some parameter regions, agents find it optimal to use their choice of a (misspecified) AR(1) rule. [Gelain, Iskrev, Lansing, and Mendicino \(2019\)](#) investigate hybrid expectations in the [Smets and Wouters \(2007\)](#) model, where some agents use moving average rules. [Hommes and Lustenhouwer \(2019\)](#) consider a NK model under heterogeneous expectations, with fundamentalists who believe in the target of the central bank versus agents with naive expectations who believe in a random walk. Along similar lines, some studies investigate *ARIMA* type forecasting rules in an experimental setup with human subjects and find evidence of small forecasting rules; see, for example, [Adam \(2007\)](#), [Beshears, Choi, Fuster, Laibson, and Madrian \(2013\)](#), and [Assenza, Heemeijer, Hommes, and Massaro \(2021\)](#).

There is much literature on optimal monetary policy rules when agents are learning. [Evans and Honkapohja \(2003, 2006\)](#) analyze the effects of learning on stability when monetary policy is conducted according to optimal policy rules under discretion and commitment and show that forward looking rules, where the policymaker observes and incorporates agents' expectations, can solve the problem of instability due to learning.⁶ [Orphanides and Williams \(2005\)](#) show that adaptive learning increases inflation persistence, which warrants a stronger policy response to inflation in order to mitigate the effects. Along similar lines, [Preston \(2006\)](#) reports that when monetary policy responds to private agents' learning behavior and decision rules, instability problems associated

⁶In their seminal paper, [Bullard and Mitra \(2002\)](#) examine the stability of the REE under variants of the standard Taylor rule and show that even when the system displays a unique, stable equilibrium under rational expectations, the parameters of the policy rule have to be chosen appropriately to ensure stability under learning.

with learning dynamics are largely avoided. Finally, [Gaspar, Smets, and Vestin \(2010\)](#) analyze how the optimal inflation and output trade-off changes when agents learn adaptively and show that the optimal targeting rule under learning resembles the optimal rule under commitment with rational expectations. Our contribution to this discussion in the literature is that we restrict our focus on a specific interest rate rule that captures the trade-off between interest rate smoothing and inflation/output gap stabilization and analyze how this trade-off changes under learning. Contrary to the literature, we expand the loss function of the central bank with an interest rate stabilization objective. We then derive numerically the coefficients capturing the trade-off between smoothing and inflation/output gap stabilization that minimizes the loss function for various weights of the interest rate stabilization objective, both under learning and under rational expectations. We show that interest rate fluctuations are more costly under learning since the central bank has to give up on inflation and output stabilization faster as the weight on interest rate stabilization rises.

2. BLE IN A MULTIVARIATE FRAMEWORK

[Hommes and Zhu \(2014\)](#) introduced BLE in the simplest setting, a one-dimensional linear stochastic model driven by an exogenous linear stochastic AR(1) process. In this paper, we generalize BLE to n -dimensional (linear) stochastic models driven by exogenous linear stochastic AR(1) processes of multiple shocks. To ease the exposition, we initially follow the presentation in [Hommes and Zhu \(2014\)](#) but generalize their 1-dimensional model to an n -dimensional framework. In addition, most macroeconomic models include lagged state variables through features such as interest rate smoothing, habit formation in consumption, investment adjustment costs, or indexation in prices and wages. Therefore, we further extend the model adding lagged state variables.

Let the law of motion of the economy be given by the stochastic difference equation

$$\mathbf{x}_t = \mathbf{F}(\mathbf{x}_{t+1}^e, \mathbf{x}_{t-1}, \mathbf{u}_t, \mathbf{v}_t), \quad (2.1)$$

where \mathbf{x}_t is an $n \times 1$ vector of endogenous variables denoted by $[x_{1t}, x_{2t}, \dots, x_{nt}]'$ and \mathbf{x}_{t+1}^e is the expected value of \mathbf{x} at date $t + 1$. Expectations may be nonrational. The map \mathbf{F} is a continuous n -dimensional vector function, \mathbf{u}_t is a vector of exogenous stationary variables, and \mathbf{v}_t is a vector of white noise disturbances.

Agents are boundedly rational and do not know the exact form of the actual law of motion (2.1). They only use a simple, parsimonious forecasting model, a univariate AR(1) process for each variable to be forecasted.⁷ Thus, agents' perceived law of motion (PLM) is assumed to be the simplest VAR model with minimum parameters, that is, a restricted VAR(1) process

$$\mathbf{x}_t = \boldsymbol{\alpha} + \boldsymbol{\beta}(\mathbf{x}_{t-1} - \boldsymbol{\alpha}) + \boldsymbol{\delta}_t, \quad (2.2)$$

⁷As shown in [Enders \(2008\)](#), parameter uncertainty increases as the model becomes more complex, and hence an estimated AR(1) model may forecast a real ARMA(2,1) process better than an estimated ARMA(2,1) model. Numerous empirical studies show that overly parsimonious models with little parameter uncertainty can provide better forecasts than models consistent with the more complex actual data-generating process (e.g., [Nelson \(1972\)](#), [Stock and Watson \(2007\)](#), [Clark and West \(2007\)](#)).

where α is a vector denoted by $[\alpha_1, \alpha_2, \dots, \alpha_n]'$, β is a diagonal matrix⁸ denoted by

$$\begin{bmatrix} \beta_1 & 0 & \cdots & 0 \\ 0 & \beta_2 & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & \beta_n \end{bmatrix}$$

with $\beta_i \in (-1, 1)$, and $\{\delta_t\}$ is a white noise process; α is the unconditional mean of x_t , and β_i is the first-order autocorrelation coefficient of variable x_i . Given the perceived law of motion (2.2), the 2-period ahead forecasting rule for x_{t+1} that minimizes the mean-squared forecasting error is

$$x_{t+1}^e = \alpha + \beta^2(x_{t-1} - \alpha). \quad (2.3)$$

Combining the expectations (2.3) and the law of motion of the economy (2.1), we obtain the implied actual law of motion (ALM),

$$x_t = F(\alpha + \beta^2(x_{t-1} - \alpha), x_{t-1}, u_t, v_t). \quad (2.4)$$

In the case where the ALM (2.4) is stationary, let the variance-covariance matrix $\Gamma(0) := E[(x_t - \bar{x})(x_t - \bar{x})']$ and the first-order autocovariance matrix $\Gamma(1) := E[(x_t - \bar{x})(x_{t+1} - \bar{x})']$, where \bar{x} is the mean of x_t . Let Ω be the diagonal matrix in which the i th diagonal element is the variance of the i th process, that is, $\Omega = \text{diag}[\gamma_{11}(0), \gamma_{22}(0), \dots, \gamma_{nn}(0)]$, where $\gamma_{ii}(0)$ is the i th diagonal entry of $\Gamma(0)$. Let L be the diagonal matrix in which the i th diagonal element is the first-order autocovariance of the i th process, that is, $L = \text{diag}[\gamma_{11}(1), \gamma_{22}(1), \dots, \gamma_{nn}(1)]$, where $\gamma_{ii}(1)$ is the i th diagonal entry of $\Gamma(1)$. Let G denote the diagonal matrix in which the i th diagonal element is the first-order autocorrelation coefficient of the i th process $x_{i,t}$. Hence,

$$G = L\Omega^{-1}. \quad (2.5)$$

Behavioral learning equilibrium (BLE) Extending on Hommes and Zhu (2014) and using the definitions of coefficients and matrices above, the concept of BLE is generalized as follows.

DEFINITION 2.1. A vector (μ, α, β) where μ is a probability measure, α is a vector, and β is a diagonal matrix with $\beta_i \in (-1, 1)$ ($i = 1, 2, \dots, n$) is called a Behavioral Learning Equilibrium (BLE) if the following three conditions are satisfied:

- S1 The probability measure μ is a nondegenerate invariant measure for the stochastic difference equation (2.4);

⁸Chung and Xiao (2013) also argue that the simple AR(1) model is more likely to prevail in reality because agents typically have restricted knowledge about the underlying system. In addition, short-term forecasts based on an AR(1) model are often better than more general VAR models because in more general VAR models too many parameters need to be estimated. Hence, coefficient uncertainty increases, leading to a deterioration in forecasting performance.

- S2 The stationary stochastic process defined by (2.4) with the invariant measure μ has an unconditional mean α , that is, the unconditional mean of x_i is α_i ($i = 1, 2, \dots, n$);
- S3 Each element x_i for the stationary stochastic process of x defined by (2.4) with the invariant measure μ has the unconditional first-order autocorrelation coefficient β_i ($i = 1, 2, \dots, n$), that is, $G = \beta$, with G defined in (2.5).

In other words, a BLE is characterized by two natural observable consistency requirements: the unconditional means and the unconditional first-order autocorrelation coefficients generated by the actual (unknown) stochastic process (2.4) coincide with the corresponding statistics for the perceived linear VAR(1) process (2.2), as given by the parameters α and β . This means that in a BLE, agents correctly perceive the two simplest and most important statistics, the mean and first-order autocorrelation (i.e., persistence) of each relevant variable of the economy, without fully understanding its structure and recognizing all explanatory variables and cross-correlations. A BLE is *parameter-free*, as the two parameters of each linear forecasting rule are pinned down by simple and observable statistics. Hence, agents do not fully understand the (linear) structure of the stochastic economy, that is, they do not observe the shocks and do not take the cross-correlations of state variables into account. Rather they use a parsimonious, but optimal univariate AR(1) forecasting rule for each state variable. A simple BLE may be a plausible outcome of the coordination process of expectations of a large population.⁹

Furthermore, along a BLE the orthogonality condition

$$E[x_{i,t} - \alpha_i - \beta_i(x_{i,t-1} - \alpha_i)] = 0,$$

$$E\{[x_{i,t} - \alpha_i - \beta_i(x_{i,t-1} - \alpha_i)]x_{i,t-1}\} = E\{[x_{i,t} - \alpha_i - \beta_i(x_{i,t-1} - \alpha_i)](x_{i,t-1} - \alpha_i)\} = 0$$

is satisfied. That is, the forecast $\alpha_i + \beta_i(x_{i,t-1} - \alpha_i)$ is the linear projection of $x_{i,t}$ on the vector $(1, x_{i,t-1})'$. For each variable, agents cannot detect the correlation between the forecasting error $x_{i,t} - \alpha_i - \beta_i(x_{i,t-1} - \alpha_i)$ and the vector $(1, x_{i,t-1})'$ in the forecast model. The linear projection produces the smallest mean-squared error among the class of linear forecasting rules (e.g., [Hamilton \(1994\)](#)). Therefore, for each variable, agents use the *optimal* forecast within their class of univariate AR(1) forecasting rules ([Branch \(2004\)](#)).

Notice that BLE is introduced by taking as given the law of motion of the economy. In other words, we do not derive the microfoundations of the model under BLE assumptions, but rather take the REE-consistent law of motion as given and introduce the new equilibrium concept. In principle, resolving the microfoundations could generate differences in the law of motion. In this paper, we abstract away from these considerations.¹⁰

⁹Laboratory experiments within the NK framework provide empirical support of the use of simple univariate AR(1) forecasting rules to forecast inflation and output gap ([Adam \(2007\)](#), [Pfajfar and Žakelj \(2014\)](#), [Assenza et al. \(2021\)](#)). See also [Hommes \(2021\)](#) for a recent survey of laboratory evidence for simple forecasting heuristics such as AR(1) rules. In Section 3.4, we will see that BLE also fits well with SPF data.

¹⁰For example, [Hommes and Zhu \(2014\)](#) solve for the microfoundations of a simple New Keynesian model and find that the reduced-form equation has a slightly different functional form; see the Online Appendix, Section 3.

Sample autocorrelation learning In the above definition of BLE, agents' beliefs are described by the linear forecasting rule (2.3) with parameters α and β fixed at their optimal values. However, the parameters α and β are usually unknown to agents. In the adaptive learning literature, it is common to assume that agents behave like econometricians using time series observations to estimate the parameters as new observations become available. Following Hommes and Sorger (1998), we assume that agents use sample autocorrelation learning (SAC-learning) to learn the parameters α_i and β_i , $i = 1, 2, \dots, n$. That is, for any finite set of observations $\{x_{i,0}, x_{i,1}, \dots, x_{i,t}\}$, the sample average is given by

$$\alpha_{i,t} = \frac{1}{t+1} \sum_{k=0}^t x_{i,k}, \quad (2.6)$$

and the first-order sample autocorrelation coefficient is given by

$$\beta_{i,t} = \frac{\sum_{k=0}^{t-1} (x_{i,k} - \alpha_{i,t})(x_{i,k+1} - \alpha_{i,t})}{\sum_{k=0}^t (x_{i,k} - \alpha_{i,t})^2}. \quad (2.7)$$

Hence, $\alpha_{i,t}$ and $\beta_{i,t}$ are updated over time as new information arrives. It is easy to check that independently of the choice of the initial values $(x_{i,0}, \alpha_{i,0}, \beta_{i,0})$, it always holds that $\beta_{i,1} = -\frac{1}{2}$ and that the first-order sample autocorrelation $\beta_{i,t} \in [-1, 1]$ for all $t \geq 1$. Similar to Hommes and Zhu (2014), we define

$$R_{i,t} = \frac{1}{t+1} \sum_{k=0}^t (x_{i,k} - \alpha_{i,t})^2.$$

Then SAC-learning is equivalent to the following recursive dynamical system:¹¹

$$\begin{cases} \alpha_{i,t} = \alpha_{i,t-1} + \frac{1}{t+1} (x_{i,t} - \alpha_{i,t-1}), \\ \beta_{i,t} = \beta_{i,t-1} + \frac{1}{t+1} R_{i,t}^{-1} \left[(x_{i,t} - \alpha_{i,t-1}) \left(x_{i,t-1} + \frac{x_{i,0}}{t+1} - \frac{t^2 + 3t + 1}{(t+1)^2} \alpha_{i,t-1} \right. \right. \\ \quad \left. \left. - \frac{1}{(t+1)^2} x_{i,t} \right) - \frac{t}{t+1} \beta_{i,t-1} (x_{i,t} - \alpha_{i,t-1})^2 \right], \\ R_{i,t} = R_{i,t-1} + \frac{1}{t+1} \left[\frac{t}{t+1} (x_{i,t} - \alpha_{i,t-1})^2 - R_{i,t-1} \right]. \end{cases} \quad (2.8)$$

The actual law of motion under SAC-learning is therefore given by

$$\mathbf{x}_t = \mathbf{F}(\alpha_{t-1} + \beta_{t-1}^2 (\mathbf{x}_{t-1} - \alpha_{t-1}), \mathbf{x}_{t-1}, \mathbf{u}_t, \mathbf{v}_t), \quad (2.9)$$

¹¹The system in (2.8) is a decreasing gain algorithm, where all observations receive equal weight and, therefore, the weight of the latest observation decreases as the sample size grows. There is also a constant gain correspondence of SAC-learning, where past observations are discounted at a geometric rate. This can be obtained by replacing the weights $\frac{1}{t+1}$ by some (small) positive constant κ . See the Online Appendix to Hommes and Zhu (2014) for further details.

with $\alpha_{i,t}$, $\beta_{i,t}$ as in (2.8). In Hommes and Zhu (2014), F is a one-dimensional linear function. In this paper, F may be an n -dimensional linear vector function and includes the lagged term \mathbf{x}_{t-1} .

2.1 Main results in a multivariate linear framework

Assume that a reduced form model is an n -dimensional linear stochastic process \mathbf{x}_t driven by an exogenous VAR(1) process \mathbf{u}_t . More precisely, the actual law of motion of the economy is given by the linear system

$$\mathbf{x}_t = F(\mathbf{x}_{t+1}^e, \mathbf{x}_{t-1}, \mathbf{u}_t, \mathbf{v}_t) = \mathbf{b}_0 + \mathbf{b}_1 \mathbf{x}_{t+1}^e + \mathbf{b}_2 \mathbf{x}_{t-1} + \mathbf{b}_3 \mathbf{u}_t + \mathbf{b}_4 \mathbf{v}_t, \quad (2.10)$$

$$\mathbf{u}_t = \mathbf{a} + \boldsymbol{\rho} \mathbf{u}_{t-1} + \boldsymbol{\varepsilon}_t, \quad (2.11)$$

where \mathbf{x}_t is an $n \times 1$ vector of endogenous variables, \mathbf{b}_0 and \mathbf{a} are vectors of constants, \mathbf{b}_1 , \mathbf{b}_2 , and \mathbf{b}_4 are $n \times n$ matrices of coefficients, \mathbf{b}_3 is an $n \times m$ matrix, $\boldsymbol{\rho}$ is an $m \times m$ matrix, \mathbf{u}_t is an $m \times 1$ vector of exogenous variables, which is assumed to follow a stationary VAR(1) as in (2.11), and \mathbf{v}_t is an $n \times 1$ vector of i.i.d. stochastic disturbance terms with mean zero and finite absolute moments and with variance-covariance matrix $\boldsymbol{\Sigma}_v$. Hence, all of the eigenvalues of $\boldsymbol{\rho}$ are assumed to be inside the unit circle. In addition, $\boldsymbol{\varepsilon}_t$ is assumed to be an $m \times 1$ vector of i.i.d. stochastic disturbance terms with mean zero and finite absolute moments. $\boldsymbol{\varepsilon}_t$ is independent of \mathbf{v}_t , and its variance-covariance matrix is $\boldsymbol{\Sigma}_\varepsilon$.

Rational expectations equilibrium

Assume that agents are rational. The perceived law of motion (PLM) corresponding to the minimum state variable REE of the model is

$$\mathbf{x}_t^* = \mathbf{c}_0 + \mathbf{c}_1 \mathbf{x}_{t-1}^* + \mathbf{c}_2 \mathbf{u}_t + \mathbf{c}_3 \mathbf{v}_t. \quad (2.12)$$

Assuming that shocks \mathbf{u}_t are observable when forecasting \mathbf{x}_{t+1} , the 1-step ahead forecast is

$$E_t \mathbf{x}_{t+1}^* = \mathbf{c}_0 + \mathbf{c}_2 \mathbf{a} + \mathbf{c}_1 \mathbf{x}_t^* + \mathbf{c}_2 \boldsymbol{\rho} \mathbf{u}_t, \quad (2.13)$$

and the corresponding actual law of motion is

$$\mathbf{x}_t^* = \mathbf{b}_0 + \mathbf{b}_1 (\mathbf{c}_0 + \mathbf{c}_2 \mathbf{a} + \mathbf{c}_1 \mathbf{x}_t^* + \mathbf{c}_2 \boldsymbol{\rho} \mathbf{u}_t) + \mathbf{b}_2 \mathbf{x}_{t-1} + \mathbf{b}_3 \mathbf{u}_t + \mathbf{b}_4 \mathbf{v}_t. \quad (2.14)$$

The REE is the fixed point of

$$\mathbf{c}_0 - \mathbf{b}_1 \mathbf{c}_1 \mathbf{c}_0 - \mathbf{b}_1 \mathbf{c}_0 = \mathbf{b}_0 + \mathbf{b}_1 \mathbf{c}_2 \mathbf{a}, \quad (2.15)$$

$$\mathbf{c}_1 - \mathbf{b}_1 \mathbf{c}_1^2 = \mathbf{b}_2, \quad (2.16)$$

$$\mathbf{c}_2 - \mathbf{b}_1 \mathbf{c}_1 \mathbf{c}_2 - \mathbf{b}_1 \mathbf{c}_2 \boldsymbol{\rho} = \mathbf{b}_3, \quad (2.17)$$

$$\mathbf{c}_3 - \mathbf{b}_1 \mathbf{c}_1 \mathbf{c}_3 = \mathbf{b}_4. \quad (2.18)$$

A straightforward computation (see Appendix A.1 in the Online Supplementary Material (Hommes, Mavromatis, Özden, and Zhu (2023)) shows that the mean of the REE $\bar{\mathbf{x}}^*$ satisfies

$$\bar{\mathbf{x}}^* = (\mathbf{I} - \mathbf{b}_1 - \mathbf{b}_2)^{-1} [\mathbf{b}_0 + \mathbf{b}_3(\mathbf{I} - \rho)^{-1} \mathbf{a}], \quad (2.19)$$

where \mathbf{I} denotes a conformable identity matrix throughout the paper. In the special case of $\rho = \rho \mathbf{I}$ and $\mathbf{b}_2 = \mathbf{0}$, the rational expectations equilibrium \mathbf{x}_t^* satisfies¹²

$$\mathbf{x}_t^* = (\mathbf{I} - \mathbf{b}_1)^{-1} \mathbf{b}_0 + (\mathbf{I} - \mathbf{b}_1)^{-1} \mathbf{b}_1 (\mathbf{I} - \rho \mathbf{b}_1)^{-1} \mathbf{b}_3 \mathbf{a} + (\mathbf{I} - \rho \mathbf{b}_1)^{-1} \mathbf{b}_3 \mathbf{u}_t + \mathbf{b}_4 \mathbf{v}_t. \quad (2.20)$$

Thus, its unconditional mean is

$$\bar{\mathbf{x}}^* = E(\mathbf{x}_t^*) = (1 - \rho)^{-1} (\mathbf{I} - \mathbf{b}_1)^{-1} [\mathbf{b}_0(1 - \rho) + \mathbf{b}_3 \mathbf{a}]. \quad (2.21)$$

Its variance-covariance matrix is

$$\begin{aligned} \Sigma_{\mathbf{x}^*} &= E[(\mathbf{x}_t^* - \bar{\mathbf{x}}^*)(\mathbf{x}_t^* - \bar{\mathbf{x}}^*)'] \\ &= (1 - \rho^2)^{-1} (\mathbf{I} - \rho \mathbf{b}_1)^{-1} \mathbf{b}_3 \Sigma_{\epsilon} [(\mathbf{I} - \rho \mathbf{b}_1)^{-1} \mathbf{b}_3]' + \mathbf{b}_4 \Sigma_v \mathbf{b}_4'. \end{aligned} \quad (2.22)$$

Furthermore, the first-order autocovariance is

$$\begin{aligned} \Sigma_{\mathbf{x}^* \mathbf{x}_1^*} &= E[(\mathbf{x}_t^* - \bar{\mathbf{x}}^*)(\mathbf{x}_{t+1}^* - \bar{\mathbf{x}}^*)'] \\ &= \rho(1 - \rho^2)^{-1} (\mathbf{I} - \rho \mathbf{b}_1)^{-1} \mathbf{b}_3 \Sigma_{\epsilon} [(\mathbf{I} - \rho \mathbf{b}_1)^{-1} \mathbf{b}_3]'. \end{aligned} \quad (2.23)$$

The first-order autocorrelation of the i th-element x_i^* of \mathbf{x}^* is the i th diagonal element of matrix $\Sigma_{\mathbf{x}^* \mathbf{x}_1^*}$ divided by the corresponding i th diagonal element of matrix $\Sigma_{\mathbf{x}^*}$. Furthermore, if $\Sigma_v = \mathbf{0}$, then the first-order autocorrelation of the i th element x_i of \mathbf{x} is equal to ρ . In this case, the persistence of the i th variable x_i^* in the REE coincides exactly with the persistence of the exogenous driving force $u_{i,t}$. That is, in this case the persistence in the REE only inherits the persistence of the exogenous driving force.

Existence of BLE

Assume that agents are boundedly rational and do not recognize that the economy is driven by an exogenous VAR(1) process \mathbf{u}_t but use simple univariate AR(1) rules to forecast the state \mathbf{x}_t of the economy. Given that agents' perceived law of motion is a restricted VAR(1) process as in (2.2), the actual law of motion is *linear* and given by

$$\mathbf{x}_t = \mathbf{b}_0 + \mathbf{b}_1 [\boldsymbol{\alpha} + \boldsymbol{\beta}^2 (\mathbf{x}_{t-1} - \boldsymbol{\alpha})] + \mathbf{b}_2 \mathbf{x}_{t-1} + \mathbf{b}_3 \mathbf{u}_t + \mathbf{b}_4 \mathbf{v}_t, \quad (2.24)$$

with \mathbf{u}_t given in (2.11). If all eigenvalues of $\mathbf{b}_1 \boldsymbol{\beta}^2 + \mathbf{b}_2$ for each $\beta_i \in [-1, 1]$, $1 \leq i \leq n$ lie inside the unit circle, then the system (2.24) of \mathbf{x}_t is stationary, and hence its mean $\bar{\mathbf{x}}$ and first-order autocorrelation \mathbf{G} exist.

¹²Note that ρ is a matrix while ρ is a scalar number, throughout the paper.

The mean of x_t in (2.24) is computed as

$$\bar{x} = (I - b_1\beta^2 - b_2)^{-1}[b_0 + b_1\alpha - b_1\beta^2\alpha + b_3(I - \rho)^{-1}a]. \quad (2.25)$$

Imposing the first consistency requirement of a BLE on the mean, that is, $\bar{x} = \alpha$, and solving for α yields

$$\alpha^* = (I - b_1 - b_2)^{-1}[b_0 + b_3(I - \rho)^{-1}a]. \quad (2.26)$$

Comparing this with (2.19), we conclude that in a BLE the unconditional mean α^* coincides with the REE mean. That is to say, in a BLE the state of the economy x_t fluctuates on average around its RE fundamental value x^* .

Consider the second consistency requirement of a BLE on the first-order autocorrelation coefficient matrix β of the PLM. The second consistency requirement yields

$$G(\beta) = \beta, \quad (2.27)$$

where $G = L\Omega^{-1}$, as in (2.5), and β are diagonal matrices. Since the actual law of motion in (2.24) is linear, the diagonal matrix $G(\beta)$ may be computed explicitly (see Appendix A.2). For convenience, let G_i denote the i th diagonal element of the matrix G in (2.5). Assuming that all of the eigenvalues of $b_1\beta^2 + b_2$ for each $\beta_i \in (-1, 1)$ ($i = 1, 2, \dots, n$) lie inside the unit circle, using the theory of stationary linear time series, $G_i(\beta_1, \beta_2, \dots, \beta_n) \in (-1, 1)$ and is a continuous function with respect to $(\beta_1, \beta_2, \dots, \beta_n)$ and other model parameters (see Appendix A.2).¹³ Based on Brouwer's fixed-point theorem for (G_1, G_2, \dots, G_n) , $\beta^* = (\beta_1^*, \beta_2^*, \dots, \beta_n^*)$ exists with each $\beta_i^* \in [-1, 1]$, such that $G(\beta^*) = \beta^*$. We conclude with the following.¹⁴

PROPOSITION 1. *If all eigenvalues of $b_1\beta^2 + b_2$ for each $\beta_i \in [-1, 1]$ are inside the unit circle, at least one behavioral learning equilibrium (α^*, β^*) exists for the economic system (2.24) with $\alpha^* = (I - b_1 - b_2)^{-1}[b_0 + b_3(I - \rho)^{-1}a] = \bar{x}^*$.*

Stability under SAC-learning

Next, we study the stability of BLE under SAC-learning. The ALM of the economy under SAC-learning is given by

$$\begin{cases} x_t = b_0 + b_1[\alpha_{t-1} + \beta_{t-1}^2(x_{t-1} - \alpha_{t-1})] + b_2x_{t-1} + b_3u_t + b_4v_t, \\ u_t = a + \rho u_{t-1} + \varepsilon_t, \end{cases} \quad (2.28)$$

¹³For example, refer to the expression (3.9) in Hommes and Zhu (2014) for the special 1-dimensional case $n = 1$ and $b_2 = 0$. In Section 2.2, we consider the NK model with two forward-looking variables, and in Appendix A.5 we compute the (complicated) expressions of $G_1(\beta_1, \beta_2)$ and $G_2(\beta_1, \beta_2)$ explicitly.

¹⁴The Schur–Cohn criterion theorem provides necessary and sufficient conditions for all eigenvalues to lie inside the unit circle (see Elaydi (2005)). For specific models, one may find sufficient conditions that are independent of β to guarantee that all eigenvalues of $b_1\beta^2 + b_2$, for each $\beta_i \in [-1, 1]$, are inside the unit circle. For example, in the case of the NK model, the Taylor principle is a sufficient condition to ensure that all eigenvalues of $b_1\beta^2 + b_2$ lie inside the unit circle for all $\beta_i \in [-1, 1]$ (see Section 2.2.2, Corollary 1, and Appendix A.4).

with α_t, β_t updated based on the realized sample average and sample autocorrelation as in (2.8). Appendix A.3 shows that the E-stability principle applies and that stability under SAC-learning is determined by the associated ordinary differential equation (ODE):¹⁵

$$\begin{cases} \frac{d\alpha}{d\tau} = \bar{x}(\alpha, \beta) - \alpha \\ \quad = (I - b_1\beta^2 - b_2)^{-1} [b_0 + b_1\alpha - b_1\beta^2\alpha + b_3(I - \rho)^{-1}a] - \alpha, \\ \frac{d\beta}{d\tau} = G(\beta) - \beta, \end{cases} \quad (2.29)$$

where $\bar{x}(\alpha, \beta)$ is the mean given by (2.25) and $G(\beta)$ is the diagonal first-order autocorrelation matrix. A BLE (α^*, β^*) corresponds to a fixed point of the ODE (2.29). Moreover, a BLE (α^*, β^*) is locally stable under SAC-learning if it is a stable fixed point of the ODE (2.29). Therefore, we have the following property of SAC-learning stability.

PROPOSITION 2. *A BLE (α^*, β^*) is locally stable (E-stable) under SAC-learning if:*

- (i) *all eigenvalues of $(I - b_1\beta^{*2} - b_2)^{-1}(b_1 + b_2 - I)$ have negative real parts, and*
- (ii) *all eigenvalues of $DG_\beta(\beta^*)$ have real parts less than 1, where DG_β is the Jacobian matrix with the (i, j) -th entry $\frac{\partial G_i}{\partial \beta_j}$.*

PROOF. See Appendix A.3.¹⁶ □

Recall from the discussion above that $G_i(\beta_1, \beta_2, \dots, \beta_n) \in (-1, 1)$, so that at least one BLE exists. Proposition 2 states when the BLE is E-stable under SAC-learning.

2.2 Application of BLE in the baseline NK model

In this section, before considering an empirical assessment of BLE, we apply our results within the framework of a standard NK model along the lines of [Gali \(2008\)](#) and [Woodford \(2003a\)](#), in order to provide an analytical comparison between BLE and REE. Consider a simple version without price indexation and habit persistence linearized around the zero inflation steady state, given by

$$\begin{cases} y_t = y_{t+1}^e - \varphi(r_t - \pi_{t+1}^e) + u_{y,t}, \\ \pi_t = \lambda \pi_{t+1}^e + \gamma y_t + u_{\pi,t}, \end{cases} \quad (2.30)$$

where y_t is the output gap, π_t is the inflation rate, and y_{t+1}^e and π_{t+1}^e are expected output gap and expected inflation, respectively. The absence of lagged state variables allows us to derive some analytical results in order to compare the BLE to the REE in this

¹⁵See [Evans and Honkapohja \(2001\)](#) for a discussion and mathematical treatment of E-stability.

¹⁶The Routh–Hurwitz criterion theorem provides sufficient and necessary conditions for all the n eigenvalues having negative real parts (see [Brock and Malliaris \(1989\)](#)).

framework. The terms $u_{y,t}$, $u_{\pi,t}$ are stochastic shocks and are assumed to follow AR(1) processes

$$u_{y,t} = \rho_y u_{y,t-1} + \varepsilon_{y,t}, \quad (2.31)$$

$$u_{\pi,t} = \rho_\pi u_{\pi,t-1} + \varepsilon_{\pi,t}, \quad (2.32)$$

where $\rho_i \in [0, 1)$ and $\{\varepsilon_{i,t}\}$ ($i = y, \pi$) are two uncorrelated i.i.d. stochastic processes with zero mean and finite absolute moments with corresponding variances σ_i^2 .

The first equation in (2.30) is an IS curve that describes the demand side of the economy. In an economy of rational or boundedly rational agents, it is a linear approximation of a representative agent's Euler equation. The parameter $\varphi > 0$ is related to the elasticity of intertemporal substitution in the consumption of a representative household, while its inverse denotes relative risk aversion. The second equation in (2.30) is the NKPC, which describes the aggregate supply relation. This is obtained by averaging all firms' optimal pricing decisions. The parameter γ is related to the degree of price stickiness in the economy, and the parameter $\lambda \in [0, 1)$ is the subjective discount factor of the representative household.

We supplement the equations in (2.30) with a standard Taylor-type policy rule, which represents the behavior of the monetary authority in setting the nominal interest rate:

$$r_t = \phi_\pi \pi_t + \phi_y y_t, \quad (2.33)$$

where r_t is the deviation of the nominal interest rate from the value that is consistent with inflation at target and output at the steady state. The parameters ϕ_π , ϕ_y , measuring the response of r_t to the deviation of inflation and output from long run steady states, are assumed to be nonnegative.

Substituting the Taylor-type policy rule (2.33) for (2.30) and writing the model in matrix form gives

$$\begin{cases} \mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1}^e + \mathbf{C}\mathbf{u}_t, \\ \mathbf{u}_t = \boldsymbol{\rho}\mathbf{u}_{t-1} + \boldsymbol{\varepsilon}_t, \end{cases} \quad (2.34)$$

where $\mathbf{x}_t = [y_t, \pi_t]'$, $\mathbf{u}_t = [u_{y,t}, u_{\pi,t}]'$, $\boldsymbol{\varepsilon}_t = [\varepsilon_{y,t}, \varepsilon_{\pi,t}]'$, $\mathbf{B} = \frac{1}{1+\gamma\varphi\phi_\pi+\varphi\phi_y} \begin{bmatrix} 1 & \varphi(1-\lambda\phi_\pi) \\ \gamma & \gamma\varphi+\lambda(1+\varphi\phi_y) \end{bmatrix}$, $\mathbf{C} = \frac{1}{1+\gamma\varphi\phi_\pi+\varphi\phi_y} \begin{bmatrix} 1 & -\varphi\phi_\pi \\ \gamma & 1+\varphi\phi_y \end{bmatrix}$, $\boldsymbol{\rho} = \begin{bmatrix} \rho_y & 0 \\ 0 & \rho_\pi \end{bmatrix}$.

Before turning to BLE, we first consider the Rational Expectations Equilibrium (REE).

2.2.1 Rational expectations equilibrium Comparing the NK model (2.34) with the general framework summarized by (2.10) and (2.11), we note that $\mathbf{a} = \mathbf{0}$, $\mathbf{b}_0 = \mathbf{0}$, and $\mathbf{b}_2 = \mathbf{0}$. The REE fixed point in (2.15)–(2.18) is then simplified to

$$(\mathbf{I} - \mathbf{B})\boldsymbol{\xi} = \mathbf{0}, \quad (2.35)$$

$$\boldsymbol{\eta} = \mathbf{B}\boldsymbol{\eta}\boldsymbol{\rho} + \mathbf{C}. \quad (2.36)$$

Bullard and Mitra (2002) show that the REE is unique (determinate) if and only if $\gamma(\phi_\pi - 1) + (1 - \lambda)\phi_y > 0$. The REE is then the stable stationary process with mean

$$\bar{\mathbf{x}}^* = \mathbf{0}. \quad (2.37)$$

In the symmetric case $\rho_i = \rho$ for $i = \{y, \pi\}$, the REE \mathbf{x}_t^* satisfies

$$\mathbf{x}_t^* = (\mathbf{I} - \rho \mathbf{B})^{-1} \mathbf{C} \mathbf{u}_t. \quad (2.38)$$

Thus, its covariance is

$$\Sigma_{\mathbf{x}^*} = \mathbf{E}(\mathbf{x}_t^* - \bar{\mathbf{x}}^*)(\mathbf{x}_t^* - \bar{\mathbf{x}}^*)' = (1 - \rho^2)^{-1} (\mathbf{I} - \rho \mathbf{B})^{-1} \mathbf{C} \Sigma_{\varepsilon} [(\mathbf{I} - \rho \mathbf{B})^{-1} \mathbf{C}]'. \quad (2.39)$$

Furthermore, the first-order autocorrelation of the i -element x_i of \mathbf{x} is equal to ρ . That is, in this case the persistence of the REE coincides exactly with the persistence of the exogenous driving force \mathbf{u}_t , and the first-order autocorrelations of output gap and inflation are the same, that is, symmetric, equal to the autocorrelation in the driving force. Therefore, in the baseline NK model without habits in consumption and price indexation, inflation, and output gap inherit the persistence of the shocks under RE.

2.2.2 Behavioral learning equilibrium As in the general setup in Section 2, we assume that agents are boundedly rational and use simple univariate linear rules to forecast the output gap y_t and inflation π_t of the economy. Therefore, we deviate from Bullard and Mitra (2002) in two important ways: (i) our agents cannot observe or do not use the exogenous shocks \mathbf{u}_t , and (ii) agents do not fully understand the linear stochastic structure and do not take into account the cross-correlation between inflation and output. Rather, our agents learn simple univariate AR(1) forecasting rules for inflation and output gap, as in (2.2). However, these AR(1) rules indirectly, in a boundedly rational way, take exogenous shocks and cross-correlations of endogenous variables into account as agents learn the two parameters of each AR(1) rule consistent with the observable sample averages and first-order autocorrelations of the state variables inflation and output gap.¹⁷

The actual law of motion (2.34) becomes

$$\begin{cases} x_t = \mathbf{B}[\boldsymbol{\alpha} + \boldsymbol{\beta}^2(x_{t-1} - \boldsymbol{\alpha})] + \mathbf{C} \mathbf{u}_t, \\ \mathbf{u}_t = \rho \mathbf{u}_{t-1} + \boldsymbol{\varepsilon}_t. \end{cases} \quad (2.40)$$

For the actual law of motion (ALM) (2.40), the REE determinacy condition $\gamma(\phi_\pi - 1) + (1 - \lambda)\phi_y > 0$ implies that the ALM is stationary for all $\boldsymbol{\beta}$ (see Appendix A.4). Thus, the means and first-order autocorrelations are

$$\begin{aligned} \bar{\mathbf{x}} &= (\mathbf{I} - \mathbf{B}\boldsymbol{\beta}^2)^{-1} (\mathbf{B}\boldsymbol{\alpha} - \mathbf{B}\boldsymbol{\beta}^2\boldsymbol{\alpha}), \\ G(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= \begin{bmatrix} G_1(\beta_y, \beta_\pi) & 0 \\ 0 & G_2(\beta_y, \beta_\pi) \end{bmatrix} = \begin{bmatrix} \text{corr}(y_t, y_{t-1}) & 0 \\ 0 & \text{corr}(\pi_t, \pi_{t-1}) \end{bmatrix}. \end{aligned}$$

For the NK model in this section without any lagged state variables, focusing on the symmetric case with $\rho_y = \rho_\pi = \rho$, we can obtain expressions for $G_1(\beta_y, \beta_\pi)$ and $G_2(\beta_y, \beta_\pi)$, which are provided in Appendix A.5. The resulting expressions depend on eight parameters φ , λ , γ , ϕ_y , ϕ_π , ρ , σ_π^2 , and σ_y^2 . Having analytical expressions for

¹⁷The use of a simple AR(1) rule is supported by evidence from the learning-to-forecast laboratory experiments in the NK framework in Adam (2007), Pfajfar and Žakelj (2014), and Assenza et al. (2021).

$G_1(\beta_y, \beta_\pi)$ and $G_2(\beta_y, \beta_\pi)$ allows us to narrow down the existence and stability conditions in this special case. Hence, using Proposition 1 and Proposition 2 we have the following properties for the NK model.

COROLLARY 1. *Under the Taylor rule (2.33), if $\gamma(\phi_\pi - 1) + (1 - \lambda)\phi_y > 0$, then at least one BLE (α^*, β^*) exists, where $\alpha^* = \mathbf{0} = \bar{x}^*$.*

COROLLARY 2. *Under the Taylor rule (2.33) and the condition $\gamma(\phi_\pi - 1) + (1 - \lambda)\phi_y > 0$, a BLE (α^*, β^*) is locally stable under SAC-learning if all eigenvalues of $DG_\beta(\beta^*) = (\frac{\partial G_i}{\partial \beta_j})_{\beta=\beta^*}$ have real parts less than 1.*

PROOF. See Appendix A.6. □

These results serve as a useful starting point to discuss some properties of BLE in a baseline setup. For the general n -dimensional case, we rely on a numerical algorithm to approximate a BLE, which is explained in Section 2.3.

To illustrate the typical output-inflation dynamics under BLE, we present a calibration exercise for empirically plausible parameter values. As in the Clarida, Gali, and Gertler (1999) calibration, we fix $\varphi = 1$, $\lambda = 0.99$. We fix $\gamma = 0.04$, which lies between the calibrations $\gamma = 0.3$ in Clarida, Gali, and Gertler (1999) and $\gamma = 0.024$ in Woodford (2003a). For the exogenous shocks, we set the ratio of shocks $\frac{\sigma_\pi}{\sigma_y} = 0.5$, which is within the possible range suggested in Fuhrer (2006). We consider the symmetric case $\rho_y = \rho_\pi = \rho = 0.5$, with weak persistence in the shocks. The baseline parameters on the policy response to inflation deviation and output gap are in line with much of the literature, $\phi_\pi = 1.5$, $\phi_y = 0.5$ (see, e.g., Fuhrer (2006, 2010)). At these parameter values, the two eigenvalues of the Jacobian matrix $DG_\beta(\beta^*)$ are $0.5012 \pm 0.7348i$ (with real parts less than 1), which implies that the BLE is E-stable under SAC-learning based on our theoretical results. The numerical results shown below are robust across a range of plausible parameter values.

Figure 1 illustrates the unique E-stable BLE $(\beta_y^*, \beta_\pi^*) = (0.9, 0.9592)$. In order to obtain (β_y^*, β_π^*) , we numerically compute the corresponding fixed-point $\beta_\pi^*(\beta_y)$, satisfying $G_2(\beta_y, \beta_\pi^*) = \beta_\pi^*$ for each β_y , and the corresponding fixed-point $\beta_y^*(\beta_\pi)$, satisfying $G_1(\beta_y^*, \beta_\pi) = \beta_y^*$ for each β_π , as illustrated in Figure 1. Hence, their intersection point (β_y^*, β_π^*) satisfies $G_1(\beta_y^*, \beta_\pi^*) = \beta_y^*$ and $G_2(\beta_y^*, \beta_\pi^*) = \beta_\pi^*$.

A striking feature of the BLE in this setup is that the first-order autocorrelation coefficients of output gap and inflation $(\beta_y^*, \beta_\pi^*) = (0.9, 0.9592)$ are substantially higher than those at the REE, that is, the persistence is much higher than the persistence $\rho (= 0.5)$ of the exogenous shocks. We refer to this phenomenon as *persistence amplification*. Agents fail to recognize the exact linear structure and cross-correlations of the economy but rather learn to coordinate the mean and the first-order autocorrelations of inflation and output gap on simple univariate AR(1) rules consistent with simple observable statistics. As a result of this *self-fulfilling mistake*, shocks to the economy are strongly amplified.

Figure 2 illustrates how these results depend on the persistence ρ of the exogenous shocks. The figure shows the BLE, that is, the first-order autocorrelations β_y^* of the output gap and β_π^* of inflation, as a function of the parameter ρ . This figure clearly shows

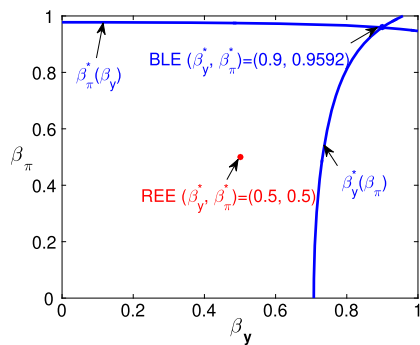


FIGURE 1. A unique BLE $(\beta_y^*, \beta_\pi^*) = (0.9, 0.9592)$ obtained as the intersection point of the fixed-point curves $\beta_\pi^*(\beta_y)$ and $\beta_y^*(\beta_\pi)$. The BLE exhibits strong persistence amplification compared to REE (red dot, with $\rho = 0.5$). Parameters are: $\lambda = 0.99$, $\varphi = 1$, $\gamma = 0.04$, $\rho = 0.5$, $\phi_\pi = 1.5$, $\phi_y = 0.5$, and $\frac{\sigma_\pi}{\sigma_y} = 0.5$.

the *persistence amplification* along BLE, with much higher persistence than under RE, for all values of $0 < \rho < 1$. Especially for $\rho \geq 0.5$, we have $\beta_y^*, \beta_\pi^* \geq 0.9$, implying that the output gap and inflation have significantly higher persistence than the exogenous driving forces. Figure 2 (right plot) also illustrates the *volatility amplification* under BLE compared to REE. For the output gap, the ratio of variances $\sigma_{y,BLE}^{*2}/\sigma_{y,REE}^{*2}$ reaches a peak of about 2.5 for $\rho \approx 0.75$, while for inflation the ratio of variances $\sigma_{\pi,BLE}^{*2}/\sigma_{\pi,REE}^{*2}$ reaches its peak at about 3.5 for $\rho \approx 0.65$.

2.3 How to find an E-stable BLE

This section discusses how to approximate a BLE. The perceived mean values α^* of a BLE are characterized by the same unconditional means as the underlying REE. Therefore, without loss of generality we may assume $\alpha^* = 0$. The first-order autocorrelation

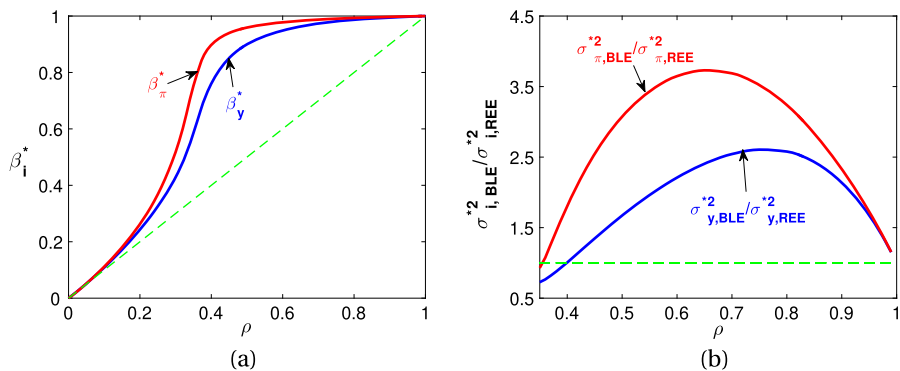


FIGURE 2. BLE (β_y^*, β_π^*) as a function of the persistence ρ of the exogenous shocks. (a) β_i^* ($i = y, \pi$) with respect to ρ ; (b) the ratio of variances $(\sigma_{y,BLE}^{*2}/\sigma_{y,REE}^{*2}, \sigma_{\pi,BLE}^{*2}/\sigma_{\pi,REE}^{*2})$ of the BLE (β_y^*, β_π^*) w.r.t. the REE. Parameters are: $\lambda = 0.99$, $\varphi = 1$, $\gamma = 0.04$, $\phi_\pi = 1.5$, $\phi_y = 0.5$, $\frac{\sigma_\pi}{\sigma_y} = 0.5$.

coefficients β^* in a BLE are functions in terms of the structural parameters μ , which satisfy the nonlinear equilibrium conditions $G(\beta^*, \mu) = \beta^*$ in (2.27), without a closed-form solution. In this section, we use the concept of *Iterative E-stability* (Evans (1985)) to find E-stable BLE for a given set of structural parameters μ .

Iterative E-stability is a simple fixed-point iteration to evaluate the mapping from perceived first-order autocorrelations β to the actual first-order autocorrelations $G(\beta, \mu)$. Given some initial conditions $\beta^{(1)}$, the iteration works as follows:

$$\beta^{(k+1)} = G(\beta^{(k)}, \mu), \quad 1 \leq k \leq N, \quad (2.41)$$

where k denotes the current iteration index, N is the total number of iterations, and μ denotes the vector of structural parameters. A BLE $(0, \beta^*)$ is locally stable under (2.41) if all eigenvalues of $DG_\beta(\beta^*)$ lie inside the unit circle. This is known as the *iterative E-stability* condition. There is a simple connection between *E-stability* and *iterative E-stability* of β^* : The former requires that the real parts of all eigenvalues of $DG_\beta(\beta^*)$ must be less than one. The latter requires that all eigenvalues of $DG_\beta(\beta^*)$ lie inside the unit circle. It follows that iterative E-stability is a stronger condition than E-stability, which leads to the following corollary.

COROLLARY 3. *Iterative E-stability of β^* implies E-stability of β^* . Therefore, if the iteration in (2.41) converges, it converges to an E-stable BLE.*

The details of the iteration procedure are discussed in Appendix B. Other practical issues in the context of estimation such as the initial values $\beta^{(1)}$ and the number of fixed-point iterations N can also be found in Appendix B.¹⁸ An advantage of using this approach as an equilibrium approximation method is that it can only converge to E-stable equilibria, which eliminates all E-unstable equilibria without additional computational steps. As a result, a BLE that converges with (2.41) is guaranteed to be stable under learning algorithms such as constant gain recursive least squares and SAC-learning.

3. EMPIRICAL APPLICATION: THE SMETS–WOUTERS MODEL

In this section, we estimate the BLE model for the canonical Smets and Wouters (2007) NK model (henceforth referred to as SW07) and consider a horse race between BLE, REE, and a variety of constant-gain Euler-equation learning models that have been used in the literature.¹⁹

We refer to BLE and REE as *equilibrium models*, where agents' PLM coefficients are fixed at their equilibrium values: the REE is pinned down by the fixed-point conditions in (2.15)–(2.18), whereas the BLE is pinned down by the fixed-point condition in (2.27). In this respect, the main difference between REE and BLE concerns knowledge about the

¹⁸Fixed-point iteration algorithms of this type have been used as an educative learning approach in earlier literature (see, e.g., DeCanio (1979), Bray (1982), Evans (1985)).

¹⁹Alternatively, one could consider constant-gain infinite horizon learning as in Preston (2005). In this paper, we only focus on Euler-equation learning models. A comparison of Euler-equation and infinite-horizon learning can be found in Gaus and Gibbs (2018).

underlying system. In a REE, agents have perfect structural knowledge of the model. In a BLE, agents do not know the cross-correlations among the variables and do not observe the shocks but use parsimonious univariate AR(1) rules and know the correct mean and first-order autocorrelation coefficients.

Our paper aims to distinguish the long-run equilibrium effects from the transient effects of learning. Adaptive learning models deviate from equilibrium models by introducing time-varying beliefs. Rather than fixing the belief coefficients at the equilibrium values, learning models allow the agents to act like econometricians and update their belief coefficients every period as new observations become available. Below we first introduce some notation to make an explicit distinction between equilibrium models BLE and REE and adaptive learning models. We then discuss the learning models that are used in our estimation exercise.

Equilibrium models

The REE and BLE models differ in terms of equilibrium computation. Once the equilibrium is solved for, each model can be represented as a recursive linear system

$$X_t = \hat{A} + \hat{B}X_{t-1} + \hat{C}\eta_t, \quad (3.1)$$

with $X_t = [x'_t, u'_t]'$, the vector of endogenous variables and exogenous AR(1) shocks, η_t , the vector of i.i.d. shocks, \hat{B} , \hat{C} , conformable matrices in terms of structural parameters, and \hat{A} , a vector of constants. BLE and REE differ in terms of \hat{B} and \hat{C} , since they satisfy different fixed-point conditions. Derivations of the matrices for both models are provided in Appendix C.1.

Adaptive learning models

In adaptive learning models, agents act like econometricians and update the belief coefficients of their PLM in every period as new observations become available. We consider a variety of learning models:

- **SAC-learning**, as described in Section 2, is the natural learning process of a BLE model where agents use a univariate AR(1) rule for every variable and update their beliefs about the mean and persistence in every period as new observations become available. Agents' PLM and the associated 2-step ahead expectations every period are given by

$$\begin{cases} x_t = \alpha_{t-1} + \beta_{t-1}(x_{t-1} - \alpha_{t-1}), \\ E_t x_{t+1} = \alpha_{t-1} + \beta_{t-1}^2(x_{t-1} - \alpha_{t-1}), \end{cases} \quad (3.2)$$

where the coefficients α_{t-1} and β_{t-1} are updated every period using SAC-learning (2.6)–(2.7) or in recursive form (2.8).

- **AR(2)-learning** with constant gain least squares is a univariate learning rule used in Slobodyan and Wouters (2012a). Agents use the following algorithm to update their

beliefs for every forward-looking variable $x_{i,t-1}$:

$$\begin{cases} R_{i,t} = R_{i,t-1} + \gamma(Y_{i,t}Y'_{i,t} - R_{i,t-1}), \\ \theta_{i,t} = \theta_{i,t-1} + \gamma R_{i,t}^{-1} Y_{i,t}(x_{i,t} - \theta_{i,t-1} Y_{i,t}), \end{cases} \quad (3.3)$$

with $\theta_{i,t} = [\alpha_{i,t}, \beta_{1,i,t}, \beta_{2,i,t}]$, $Y_{i,t} = [1, x_{i,t-1}, x_{i,t-2}]'$, and $R_{i,t}$ the perceived variance of the variable $x_{i,t}$.²⁰ A potential advantage of this PLM over the AR(1) rule is that it can generate an extrapolation bias in beliefs, where the most recent observation receives more weight relative to its AR(1) counterpart and the second lagged variable gets negative weight.²¹

- **Pseudo MSV-learning** with constant-gain least squares where agents use the correctly specified functional form associated with a REE, namely the MSV solution of the model, but are uncertain about its parameters. Their PLM and the associated 2-step ahead expectations at period t are given by

$$\begin{cases} \mathbf{x}_t = \gamma_{0,t-1} + \gamma_{1,t-1}\mathbf{x}_{t-2} + \gamma_{2,t-1}\mathbf{u}_{t-1}, \\ E_t\mathbf{x}_{t+1} = \gamma_{0,t-1} + \gamma_{1,t-1}\mathbf{x}_{t-1} + \gamma_{2,t-1}\rho\mathbf{u}_{t-1}, \end{cases} \quad (3.4)$$

which depends on both state variables \mathbf{x}_{t-1} and exogenous AR(1) shocks \mathbf{u}_{t-1} . Agents' learning algorithm assumes the same functional form as in (3.3) in multivariate form:

$$\begin{cases} R_t = R_{t-1} + \gamma(Y_t Y'_t - R_{t-1}), \\ \theta_t = \theta_{t-1} + \gamma R_t^{-1} Y_t(x_t - \theta_{t-1} Y_t), \end{cases} \quad (3.5)$$

where Y_t consists of a 15×1 vector of endogenous variables, exogenous shocks, and an intercept. θ_t is a 15×15 matrix of PLM coefficients.²²

- **VAR(1)-learning** with constant gain least squares where agents use only the state variables. This has been referred to as *limited information learning* (Xiao and Xu (2014)) in the literature and corresponds to a restricted version of the MSV-learning model described above. In VAR(1)-learning, agents use the following PLM and 2-step ahead expectations:

$$\begin{cases} \mathbf{x}_t = \gamma_{0,t-1} + \gamma_{1,t-1}\mathbf{x}_{t-2}, \\ E_t\mathbf{x}_{t+1} = \gamma_{0,t-1} + \gamma_{1,t-1}\mathbf{x}_{t-1}. \end{cases} \quad (3.6)$$

²⁰A generalization of the SAC-learning algorithm to other types of PLMs, such as AR(2), is undertaken in Branch, Evans, and McGough (2014). In this paper, we apply this learning method to AR(1)-learning only and use the standard constant-gain recursive least squares for other learning models.

²¹Empirical evidence in favor of such an extrapolation bias has been found in, for example, Fuster, Hebert, and Laibson (2010) and Bordalo, Gennaioli, Ma, and Shleifer (2020). An assessment of alternative theoretical approaches that support extrapolating expectations, with an initial underreaction to shocks followed by a delayed overreaction, can be found in Angeletos, Huo, and Sastry (2021).

²²This corresponds to 7 state variables, 7 exogenous shocks, and the intercept in the context of the SW07 model. The government spending shock g_t in the model is highly correlated with output y_t . Therefore, we exclude g_t from agents' regression model (3.5) when estimating the model in practice, which improves the performance of the pseudo-MSV learning model.

Agents' learning algorithm assumes the same functional form as in (3.5), where Y_t consists of an 8×1 vector of endogenous variables and an intercept. θ_t is an 8×8 matrix of PLM coefficients. This specification helps us bridge the gap between univariate AR(1)–AR(2) models and the REE-consistent knowledge. Compared to BLE, VAR(1) takes the cross-correlations into account, while BLE uses univariate AR(1) rules.

Similar to the equilibrium models, learning models can be represented as a recursive linear system after plugging in the expectations:

$$X_t = \hat{A}_{t-1} + \hat{B}_{t-1}X_{t-1} + \hat{C}_{t-1}\eta_t, \quad (3.7)$$

with time-varying matrices \hat{B}_{t-1} , \hat{C}_{t-1} and perceived mean vector \hat{A}_{t-1} , where the time variation comes from agents' PLM coefficients. Derivations of the matrices for all learning models are provided in Appendix C.2.

3.1 Estimation methodology and other practical issues

Timing of expectations and the Kalman filter Both BLE and REE equilibrium models admit a multivariate linear structure and, therefore, the likelihood function can be evaluated using standard Kalman filter recursions. For the learning models, we assume a sequential timing of intraperiod events as follows:

1. Shocks u_t are realized.
2. Expectations $E_t x_{t+1}$ are formed based on the previous period's state variables x_{t-1} , exogenous shocks u_{t-1} , and belief coefficients θ_{t-1} .
3. State variables x_t are realized.
4. Belief coefficients θ_t are updated based on period t realizations of x_t and shocks u_t .

This structure assumes that expectations and belief coefficients are predetermined before the state variables are realized. This is known as $t - 1$ **timing of expectations** in the literature. The advantage of this approach is that it allows for a conditionally linear model structure and, therefore, the likelihood function for learning models can be evaluated using the standard Kalman filter. Similar assumptions have been used elsewhere in the literature to make use of standard likelihood methods, such as Milani (2005), Slobodyan and Wouters (2012a, 2012b), and Jääskelä and McKibbin (2010). The details of the Kalman filter recursions are discussed in Appendix D.

The timing structure of expectations in our learning models differs from the t -timing of expectations that is often assumed in REE models. In a REE, expectations and state variables are jointly realized, that is, agents fully internalize period t information when forming their expectations.²³ In our paper, we abstract away from these considerations

²³Previous studies in the literature such as Milani (2005) and Slobodyan and Wouters (2012b) have used predetermined belief coefficients together with a joint determination of expectations and state variables. While this approach still admits a conditionally linear structure that can be used with a Kalman filter, it in-

and use the term *pseudo MSV-learning* to make a clear distinction between our approach and learning with fully rational knowledge about the structure of the underlying system.

Note that generally, not all forward-looking and state-variables are observed in the estimation.²⁴ In cases where agents' beliefs depend on unobserved state variables, we assume that they know the Kalman filter estimates of these variables. In other words, agents and the econometrician have fundamentally different information sets where agents are implicitly assumed to know the Kalman filter estimates when forming their expectations. This is a standard assumption when estimating DSGE models, and we do not account for the uncertainty around unobserved state variables in this paper.

Initial beliefs A practical issue when it comes to estimating adaptive learning models is the initialization of beliefs. Many studies have shown that initial beliefs matter when it comes to empirical performance of learning models, for example, [Slobodyan and Wouters \(2012b\)](#), [Berardi and Galimberti \(2017\)](#), and [Gaus and Gibbs \(2018\)](#), among others. In particular, [Gaus and Gibbs \(2018\)](#) decompose the improvements associated with learning models into two components: the role of initial beliefs and the role of time variation in beliefs. They find that within the class of PLMs that nest the MSV solution in their model, initial beliefs play a more important role in driving model fitness than the time variation in beliefs.

Our goal in this paper is not to assess the impact of initial beliefs on the performance of learning models. Rather, we are interested in using a reasonable initialization benchmark for learning models to compare against the equilibrium models BLE and REE. Therefore, we adopt a practical regression-based approach to initialize the learning models: we simulate data from our estimated BLE and REE models and run a regression to obtain initial beliefs consistent with the knowledge about the economy associated with each learning model. For SAC- and AR(2)-learning (models with univariate learning rules), we use simulated data from BLE to initialize them. For pseudo MSV- and VAR(1)-learning (models with multivariate learning rules), we use simulated data from REE. Using an underlying equilibrium concept for belief initialization is consistent with the approaches in [Slobodyan and Wouters \(2012a, 2012b\)](#). Further, [Berardi and Galimberti \(2017\)](#) suggest that equilibrium-related initialization methods result in more robust parameter estimates and are less prone to small sample size issues compared to other alternatives.

Projection facilities Another practical matter in learning models is the implementation of projection facilities. When estimating these models, some parameter and shock combinations may lead to updates in learning coefficients that imply explosive dynamics and unstable outcomes. A standard approach in learning literature is to discard the updates on learning coefficients if the new draws generate explosive dynamics (see, e.g.,

introduces a timing inconsistency for the agents: While their expectations are based on period t information, their belief coefficients are based on period $t - 1$. Therefore, we assume $t - 1$ timing on both ends for all models considered to have a consistent treatment.

²⁴For example, in the context of the SW07 model, asset prices and real interest rate of capital are unobserved, whereas consumption, investment, and real wages are only observed in growth rates.

Milani (2005) and Slobodyan and Wouters (2012b)). In this paper, we follow a similar approach and discard belief updates that generate unstable ALMs.²⁵

Model, priors, and measurement equations We use quarterly U.S. data over the period 1966:I–2007:IV to estimate the models. We repeat the estimation exercise with two sets of observable variables with and without inflation survey expectations:

- First, we follow the original Smets and Wouters (2007) model structure and use 7 observable variables: the (log-) difference of real GDP, real consumption, real investment, real wages, (log-) hours worked, CPI inflation,²⁶ and the federal funds rate.
- Second, we reestimate the models by additionally including short-term (1-quarter ahead) inflation expectations from the SPF (Croushore (1993)). This approach follows Carvalho, Eusepi, Moench, and Preston (2023), where the models are estimated using short-term inflation expectations data only.²⁷

We treat the model with the original set of observables as our baseline specification to evaluate the in-sample and pseudo out-of-sample forecasting performance of the models. In Section 3.4, we use the reestimation results with inflation expectations to discuss how the models fit survey data.

Our model follows the original Smets and Wouters (2007) structure with minor deviations (see Appendix E for further details). The model consists of 13 equations with 7 forward-looking variables, 7 exogenous AR(1) shocks, and 7 state variables. There are 35 estimated parameters including the constant gain for the adaptive learning models. We leave further details of the model, measurement equations, and the prior distributions to Appendix E.

Both equilibrium and adaptive learning models are estimated using a standard Kalman filter combined with Bayesian likelihood methods. For all models, we first obtain the posterior mode using Sims' (1999) *csmmwel* algorithm. We use the estimated posterior as a candidate density to initialize the Monte Carlo Markov Chains (MCMC), where we use a random-walk Metropolis–Hastings algorithm. For each model, we use two parallel Markov Chains where the scale coefficient of the covariance matrix is used to obtain an acceptance ratio between 30 and 45%. Each Markov Chain contains 500,000 draws, where the first half is discarded as a burn-in sample and the second half is used to compute the posterior moments and Modified Harmonic Mean (MHM) estimates. Further details of the Kalman filter and the estimation procedure for both equilibrium and learning models are outlined in Appendix D.

²⁵Note that for SAC-learning a projection facility is not needed, as the autocorrelation coefficients always lie in the interval $[-1, +1]$.

²⁶Note that Smets and Wouters (2007) use the GDP deflator as their inflation measure. We use CPI inflation in our estimations in order to make use of the survey data available in the SPF.

²⁷Carvalho et al. (2023) use their estimates to evaluate the models' performance in matching long-run inflation expectations. Here, we abstract away from a formal evaluation of long-run expectations and discuss the implications for this only qualitatively.

3.2 Baseline estimation results

Table 1 shows the posterior mean estimates for all 6 models in our baseline setup. We discuss the estimation results along two dimensions: model fitness, based on the MHM, and differences in the estimated parameter values. We introduce Bayes' factors relative to the REE benchmark in the last row of the table.²⁸

The overall pattern in model fitness suggests that the BLE model, as well as all learning models, outperforms the REE benchmark, with all Bayes' factors exceeding 4. The BLE model yields a fitness comparable to pseudo MSV- and AR(2)-learning models, while SAC- and VAR(1)-learning models generate the best outcomes in terms of model fitness.²⁹ These results suggest that (i) the knowledge about the underlying system on expectations (BLE vs. REE), in isolation from any learning effects, plays an important role in driving the model fit, and (ii) learning improves the fit, but the degree of improvements in the learning models depends on the degree of knowledge about the underlying system that the agents are using. In particular, BLE explains about 75% of the improved fit under SAC-learning (Bayes' factors 6.87 vs. 9.30).

In order to discuss differences in parameter estimates across models, we divide the parameters into four main buckets: structural parameters that determine endogenous persistence and slopes in Euler equations and Phillips curves; monetary policy parameters that appear in the Taylor rule reaction function; parameters related to steady state and measurement equations of the model; and shock persistence and standard deviations.

For monetary policy and steady-state groups, we do not observe important differences in parameter estimates across the models, and all models feature HPD intervals well within the range of each other. There are some differences in the estimated shock persistence and structural parameter groups. To understand the intuition behind these differences, we first cover the main portion of the model that interacts with expectations.³⁰ The consumption Euler equation in the model is given by

$$\begin{cases} c_t = c_1 c_{t-1} + (1 - c_1) \mathbb{E}_t c_{t+1} + c_2 (l_t - \mathbb{E}_t l_{t+1}) - c_3 (r_t - \mathbb{E}_t \pi_{t+1}) + \epsilon_t^b, \\ \epsilon_t^b = \rho_b \epsilon_{t-1}^b + \eta_t^b, \end{cases} \quad (3.8)$$

with $c_1 = \frac{\lambda}{\gamma} / (1 + \frac{\lambda}{\gamma})$, $c_2 = (\sigma_c - 1)(w_{ss} l_{ss} / c_{ss}) / (\sigma_c (1 + \frac{\lambda}{\gamma}))$, $c_3 = (1 - \frac{\lambda}{\gamma}) / ((1 + \frac{\lambda}{\gamma}) \sigma_c)$. Similarly, the investment Euler equation is given by

$$\begin{cases} i_t = i_1 i_{t-1} + (1 - i_1) \mathbb{E}_t i_{t+1} + i_2 q_t + \epsilon_t^i, \\ \epsilon_t^i = \rho_i \epsilon_{t-1}^i + \eta_t^i, \end{cases} \quad (3.9)$$

²⁸The Bayes' factors are computed as the likelihood (MHM) ratio of each model relative to REE, normalized by common logarithm base 10. We use Jeffrey's guidelines (Greenberg (2012)) to compare the Bayes' factors, which suggests that a Bayes' factor larger than 2 can be interpreted as providing *decisive support* for the model under consideration, relative to the REE benchmark.

²⁹Our results on pseudo MSV-learning are in line with previous estimates reported in Milani (2007) and Slobodyan and Wouters (2012b). The Bayes' factors implied by their results are 2.8 and 5.1, respectively. As such, our estimate of 4.72 falls within this range.

³⁰The remaining model equations can be found in Appendix E.

TABLE 1. Estimation results (posterior means) with 7 observables—no inflation expectations.

Parameter	Equilibrium Models		Learning Models			
	REE	BLE	SAC	Pseudo MSV	VAR(1)	AR(2)
Structural Parameters						
ϕ (Capital adj. cost)	5.68	2.12	1.38	5.17	2.23	2.21
σ_c (Inv. elasticity of subs.)	1.3	0.5	0.52	1.69	0.9	0.6
λ (Habit formation)	0.77	0.83	0.71	0.71	0.69	0.8
ξ_w (Wage Calvo)	0.74	0.72	0.73	0.69	0.71	0.68
σ_l (Elasticity of labor supply)	1.29	2.5	2.81	1.87	2.29	1.27
ξ_p (Price Calvo)	0.59	0.71	0.52	0.67	0.61	0.54
ι_w (Wage indexation)	0.31	0.14	0.16	0.35	0.2	0.16
ι_p (Price indexation)	0.2	0.5	0.46	0.39	0.46	0.33
ψ (Elasticity of capital util.)	0.55	0.5	0.47	0.33	0.47	0.46
ϕ_p (Production fixed costs)	1.65	1.41	1.36	1.59	1.54	1.47
α (Capital share of output)	0.17	0.14	0.13	0.18	0.16	0.15
Monetary Policy						
ϕ_π (Inflation reaction)	1.51	1.51	1.61	1.46	1.41	1.46
ρ (Smoothing)	0.86	0.91	0.91	0.91	0.92	0.9
ϕ_y (Output gap reaction)	0.11	0.11	0.14	0.13	0.11	0.11
$\phi_{\Delta y}$ (Output gap growth reaction)	0.15	0.13	0.14	0.13	0.12	0.12
Steady State						
$\bar{\pi}$ (Inflation S.S.)	0.69	0.77	0.74	0.77	0.77	0.74
$\bar{\beta}$ (Discount factor)	0.17	0.27	0.28	0.27	0.26	0.31
\bar{l} (Hours worked S.S.)	1.2	−0.12	−0.3	−0.62	−1.12	−2.04
$\bar{\gamma}$ (S.S. growth rate)	0.4	0.41	0.42	0.41	0.42	0.4
Shock Persistence						
ρ_a (TFP)	0.92	0.93	0.94	0.91	0.93	0.93
ρ_b (Risk premium)	0.34	0.32	0.46	0.19	0.18	0.4
ρ_g (Gov. spending)	0.99	0.98	0.97	0.97	0.97	0.97
ρ_i (Investment)	0.8	0.44	0.55	0.58	0.46	0.5
ρ_r (Monetary policy)	0.08	0.11	0.1	0.1	0.11	0.11
ρ_p (Price mark-up)	0.59	0.08	0.12	0.46	0.1	0.07
ρ_w (Wage mark-up)	0.84	0.3	0.38	0.86	0.13	0.25
ρ_{ga} (TFP impact on Gov.)	0.5	0.54	0.54	0.54	0.54	0.52
Shock St. Dev.						
η_a (Productivity)	0.45	0.48	0.5	0.45	0.46	0.47
η_b (Risk premium)	2.35	4.4	2.57	2.74	3.21	4.24
η_g (Gov. spending)	0.56	0.5	0.49	0.51	0.5	0.5
η_i (Investment)	0.39	1.5	1.55	1.76	1.69	1.58
η_r (Monetary policy)	0.22	0.21	0.21	0.22	0.21	0.21
η_p (Price mark-up)	0.21	0.53	0.53	0.23	0.5	0.53
η_w (Wage mark-up)	0.11	0.58	0.61	0.11	0.58	0.59
constant gain			0.006	0.008	0.024	0.008
(Log-) likl at mode	−1069.08	−1049.35	−1043.59	−1055.02	−1049.4	−1043.87
MHM	−1143.09	−1127.26	−1121.66	−1132.21	−1122.34	−1130.82
Bayes' factor	0	6.87	9.30	4.72	9.01	5.33

with $i_1 = \frac{1}{1+\bar{\beta}\gamma}$, $i_2 = \frac{1}{(1+\bar{\beta}\gamma)(\gamma^2\phi)}$, where $\bar{\beta} = \beta\gamma^{-\sigma_c}$. The price NKPC equation is

$$\begin{cases} \pi_t = \pi_1 \mathbb{E}_t \pi_{t+1} - \pi_2 \mu_t^p + \epsilon_t^p, \\ \epsilon_t^p = \rho_p \epsilon_{t-1}^p + \eta_t^p, \end{cases} \quad (3.10)$$

with $\pi_1 = \bar{\beta}\gamma$, $\pi_2 = (1 - \beta\gamma\xi_p)(1 - \xi_p)/[\xi_p((\phi_p - 1)\epsilon_p + 1)]$. The wage Phillips curve equation is

$$\begin{cases} w_t = w_1 w_{t-1} + (1 - w_1)(\mathbb{E}_t w_{t+1} + \mathbb{E}_t \pi_{t+1}) - w_2 \mu_t^w + \epsilon_t^w, \\ \epsilon_t^w = \rho_w \epsilon_{t-1}^w + \eta_t^w, \end{cases} \quad (3.11)$$

with $w_1 = 1/(1 + \bar{\beta}\gamma)$ and $w_2 = ((1 - \bar{\beta}\gamma\xi_w)(1 - \xi_w)/(\xi_w(\phi_w - 1)\epsilon_w + 1))$. Finally, the capital asset pricing equation (Tobin's q) is

$$q_t = q_1 \mathbb{E}_t q_{t+1} + (1 - q_1) \mathbb{E}_t r_{t+1}^k - (r_t - \mathbb{E}_t \pi_{t+1}) + \frac{1}{c_3} \epsilon_t^b, \quad (3.12)$$

with $q_1 = \bar{\beta}(1 - \delta)$. Among the shock persistence terms, investment shock ϵ_t^i and wage mark-up shock ϵ_t^w are more persistent under REE compared to BLE and all 4 learning models. These shocks enter the model through investment Euler equation (3.9) and the wage Phillips curve (3.11), respectively. The results suggest that both backward-looking expectations in BLE and time-varying expectations in learning models are able to capture some of the exogenous persistence in these equations through the expectation terms. The remaining shocks are comparable across all models in terms of persistence and volatility.

Among the structural parameters, capital adjustment cost ϕ and the inverse of the elasticity of intertemporal substitution σ_c stand out as the biggest differences among the models, where both parameters are smaller under the BLE and learning models compared to REE. σ_c has a two-fold effect: First, it determines the feedback from the real interest rates ($r_t - \mathbb{E}_t \pi_{t+1}$) on consumption and Tobin's q , as shown in (3.8) and (3.12), respectively. The estimated parameter is smaller in the BLE and learning models, which translates into a stronger feedback channel. Second, σ_c determines the relation between expected change in hours worked ($l_t - \mathbb{E}_t l_{t+1}$) and consumption. $\sigma_c > 1$ implies complementarity between expected change in hours worked and consumption, whereas $\sigma_c < 1$ implies that they are substitutes. The results suggest that they are complements under REE and MSV-learning, whereas they are substitutes under BLE and other learning models. The key driver for these results is how the shocks interact with expectations and model equations: in REE and pseudo MSV-learning models, the shocks enter the model equations through the expectation terms, which introduces a positive correlation between consumption and expected change in hours worked in REE and pseudo MSV-learning models. When we use an AR(1), AR(2), or VAR(1) information set instead, the mean reversion in hours worked plays a stronger role and drives the negative correlation between hours worked and consumption. For the remaining structural parameters, in particular the Calvo probabilities and indexation terms, there are no systematic differences between REE, BLE, and learning models.

Taken together, we find that both BLE and learning models improve the model fit relative to REE, without substantially affecting most parameter estimates. These results are consistent with the findings in Jääskelä and McKibbin (2010) and Slobodyan and Wouters (2012a, 2012b). Our results also complement the analysis in Gaus and Gibbs (2018), who document that initial beliefs play a more important role in driving the model fit than the time variation in beliefs within the class of PLMs that take the form of an MSV solution. We show that similar results hold for AR(1) beliefs that do not nest the MSV solution. Replacing the REE-consistent PLM with simple AR(1) beliefs (REE vs. BLE) improves the fit more than introducing time variation in AR(1) beliefs (BLE vs. SAC).

3.3 Pseudo out-of-sample forecasts

In this section, we use the 6 models presented in Table 1 and consider a pseudo out-of-sample forecasting (POOS) exercise. For each model, we use a rolling-window estimation starting with the 20-year period 1966:I–1986:IV. We reestimate the models at each quarter by rolling forward the estimation window and compute the associated out-of-sample forecast errors up to 12 quarters ahead for all observable variables. In learning models, the initial beliefs are updated every period using the same methodology as in Section 3.2. As such, we first reestimate the REE and BLE models for each period. Then we update the initial beliefs for learning models using simulated data from reestimated REE and BLE models at every period.

We compute the forecast errors associated with each model and report the *percentage changes in RMSEs relative to REE* for the BLE and learning models in Table 2. The relative RMSEs are computed as the percentage difference in RMSEs between the REE benchmark and each model: A positive (negative) number in Table 2 reflects the percentage gains (losses) in forecasting performance for the associated model relative to REE. The last column in Table 2 reports a summary statistic for each model using the *uncentered log-determinant of the forecast error covariance matrix* of all 7 observable variables.³¹

The forecasting performance of both the BLE and learning models relative to REE is characterized by an inverse U-shaped pattern: All models outperform the REE benchmark up to 4Q ahead, resulting in performance gains of up to 17%. The forecasting performance typically deteriorates at longer horizons, and the forecasts are generally worse than the REE with 88 and 12-quarter ahead forecasts. These results are consistent with the findings reported in Slobodyan and Wouters (2012b), which compare an AR(2) model with Kalman-gain learning to the REE benchmark. The results suggest that cross-restrictions imposed by the REE model are useful particularly over longer horizons, while the BLE and learning models with limited knowledge about the underlying system provide more accurate forecasts over shorter horizons.

Looking at the relative RMSEs for individual variables reveals that output, consumption, investment, and wage growth forecasts are generally comparable to or better than REE, both in the short and long run, for both the BLE and learning models, while the

³¹The summary statistic measure follows the approach in Smets and Wouters (2007).

TABLE 2. Percentage differences in RMSEs relative to the rational expectations model. A positive (negative) number reflects the percentage gains (losses) in forecasting performance relative to REE.

Horizon	Δy_t	Δc_t	Δinv_t	Δw_t	π_t	r_t	l_t	Summary
BLE								
1Q	-0.48	8.11	0.92	4.86	17.91	22.62	18.59	12.28
2Q	-2.75	19	-6.93	2.62	27.28	30.92	15.75	13.71
4Q	1.27	23.52	-1.66	1.7	34.15	29.66	3.09	17.05
8Q	10.8	23.59	2.13	-1.27	-6.15	7.51	-5.61	0.14
12Q	6.75	15.94	0.04	-6.66	-32.4	-13.06	0.95	-6.77
pseudo MSV								
1Q	-3.59	5.86	-10.7	0.12	-11.6	1.77	10.91	2.55
2Q	-5.49	12.37	-15.4	-1.46	-10.9	8.48	12.09	4.2
4Q	1.57	19.44	-5.64	-4.81	-21	3.35	5.46	6.32
8Q	9.86	15.95	3.82	0.82	-70.5	-18.63	6.48	-4.23
12Q	2.4	3.11	2.12	1.33	-91.4	-36.07	4.2	-9.07
SAC								
1Q	4.08	2.06	1.56	-2.83	21.82	19.06	17.65	9.36
2Q	3.64	9.91	-3.18	-3.32	29.16	26.28	17.37	11.53
4Q	7.79	16.72	-0.39	-0.13	33.45	22.87	9.9	14.86
8Q	12.8	18.97	2.68	0.17	4.87	2.8	9.1	2.84
12Q	5.68	8.46	1.52	-0.4	-29.1	-16.96	15.36	-4.7
pseudo-VAR(1)								
1Q	-1.24	10.72	-1.11	2.06	17.16	18.16	13.41	8.7
2Q	-2.85	15.48	-6.4	0.4	14.9	26.13	12.25	8.05
4Q	1.83	21.56	-6.65	1.09	3.74	21.61	1.75	6.88
8Q	13.3	21.48	0.49	0.92	-19.7	-2.85	-7.48	-2.43
12Q	11.2	10.82	8	0.2	-41.2	-31.27	8.78	-4.24
AR(2)								
1Q	-1.41	4.79	0.57	-5.44	10.98	16.62	15.52	7.63
2Q	-5.98	14.4	-6.51	-6.59	26.36	21.05	9.95	10.72
4Q	-2.62	20.47	-3.9	-3.98	30.92	12.42	-7.58	13.25
8Q	6.84	21.27	1.81	0.64	-0.24	-24.67	-24.23	-0.26
12Q	3.12	12.9	1.55	0.08	-35	-58.67	-17.48	-9.3

trade-off between the short and long run is driven mainly by inflation and interest rate forecasts. With the exception of the pseudo MSV model, all models outperform inflation and interest rate forecasts of REE in the short run, while they are outperformed in the long run.

An important takeaway from the POOS exercise is that the forecasting performance of the BLE model is competitive with learning models, and both BLE and learning models improve the forecasting performance relative to REE up to 4 quarters ahead. This suggests that when deviating from the REE benchmark, both the time variation in beliefs and the degree of knowledge about the underlying system imposed on the agents play an important role. In the next section, we extend the baseline estimation results reported in Table 1 to incorporate short-term inflation survey expectations.

3.4 Inflation expectations

In this section, we extend the baseline estimation results reported in Table 1 to incorporate short-term inflation expectations. In particular, we use 1-quarter ahead inflation expectations from the SPF for the U.S. For each model, we use the following identity to link the model-implied inflation expectations to the data:

$$\left\{ \pi_{t+1}^{\text{SPF}} = \mathbb{E}_t \pi_{t+1} + \eta_t^{\pi^{\text{exp}}}, \right. \quad (3.13)$$

with π_{t+1}^{SPF} referring to the SPF forecasts, $\mathbb{E}_t \pi_{t+1}$, the model-implied 2-step ahead inflation expectations, and $\eta_t^{\pi^{\text{exp}}}$, an IID measurement error. We use the same estimation period 1966:I–2007:IV. Since SPF data is only available from 1983:III onwards, we treat inflation expectations as unobserved for the earlier sample period 1966:I–1982:II.³²

Table 3 reports the estimation results and posterior means for all models. The parameter estimates are generally in line with those in Table 1, suggesting that the inclusion of short-term inflation expectations data does not lead to substantial differences in the model structure. Some notable exceptions among the structural parameters include the Calvo probabilities, price and wage indexations, and the elasticity of labor supply. These parameters interact directly with inflation expectations through the price and wage NKPCs (3.10) and (3.11), respectively. In particular, for the REE model, the wage NKPC becomes steeper (lower-wage Calvo parameter, ξ_w), while the price NKPC becomes flatter (higher-price Calvo parameter, ξ_p). The same pattern is also evident for the pseudo MSV-learning model as regards the price NKPC, while the changes in the respective parameter estimates in the BLE and the other learning models are negligible.

The Bayes' factors in Table 3 with expectations survey data are significantly larger than those in Table 1 without survey expectations: while the Bayes' factors in Table 1 without inflation expectations range between 4.72 and 9.30, the range in Table 3 increases to 35.47–53.54. This suggests that the gap in model fitness relative to the REE benchmark widens for the BLE and all learning models. The results on learning models suggest that time-varying dynamics help to capture the expectation dynamics better, which is consistent with the findings in [Carvalho et al. \(2023\)](#), [Slobodyan and Wouters \(2012a, 2012b\)](#), and [Ormeño and Molnár \(2015\)](#). A novelty of our results is that the BLE model, an equilibrium model with fixed beliefs, is competitive with learning models even after inflation expectations survey data are included as observables. BLE explains about 80% of the improved fit of SAC-learning (Bayes' factors 42.74 vs. 53.54).

To understand how well the models fit inflation expectations data, we show the model-implied inflation expectations against survey data in Figure 3 and some correlation statistics in Table 4.³³ A noticeable feature of both BLE and learning models is that they imply high inflation expectations during the 1970s and 1980s in the high

³²In this paper, we only consider an analysis of survey data on inflation expectations. Since we consider a deviation from rational expectations for all forward-looking variables in our BLE and learning models, a similar analysis can also be extended to expectations on aggregate consumption, investment, and all other forward-looking variables depending on the availability of data. We leave these considerations to future work and only focus on inflation dynamics in this paper.

³³For model-implied expectations, we refer to $\mathbb{E}_t \pi_{t+1}$ in (3.13) in the absence of any measurement errors.

TABLE 3. Estimation results (posterior means) with 8 observables, including 1-quarter ahead inflation expectations.

Parameter	Equilibrium Models		Learning Models			
	REE	BLE	SAC	Pseudo MSV	VAR(1)	AR(2)
Structural Parameters						
ϕ (Capital adj. cost)	5.04	1.36	1.84	4.78	3.37	2.49
σ_c (Inv. elasticity of subs.)	1.4	0.51	0.68	0.98	0.79	0.63
λ (Habit formation)	0.71	0.75	0.72	0.69	0.74	0.77
ξ_w (Wage Calvo)	0.45	0.73	0.68	0.63	0.68	0.73
σ_l (Elasticity of labor)	2.89	1.9	2.24	1.72	1.31	1.64
ξ_p (Price Calvo)	0.86	0.72	0.6	0.82	0.55	0.56
ι_w (Wage indexation)	0.12	0.22	0.22	0.15	0.32	0.31
ι_p (Price indexation)	0.22	0.4	0.28	0.52	0.19	0.24
ψ (Elasticity of capital util.)	0.44	0.49	0.5	0.49	0.47	0.52
ϕ_p (Production fixed costs)	1.71	1.42	1.53	1.56	1.55	1.51
α (Capital share of output)	0.2	0.14	0.16	0.17	0.16	0.15
Monetary Policy						
ϕ_π (Inflation reaction)	1.61	1.56	1.5	1.51	1.66	1.46
ρ (Smoothing)	0.85	0.9	0.9	0.9	0.9	0.89
ϕ_y (Output gap reaction)	0.11	0.11	0.13	0.08	0.13	0.12
$\phi_{\Delta y}$ (Output gap growth reaction)	0.16	0.14	0.13	0.11	0.13	0.13
Steady State						
$\bar{\pi}$ (Inflation S.S.)	0.8	0.84	0.63	0.49	0.77	0.72
$\bar{\beta}$ (Discount factor)	0.25	0.24	0.26	0.29	0.26	0.26
\bar{l} (Hours worked S.S.)	1.32	-0.52	-0.2	2.37	0.86	-1.08
$\bar{\gamma}$ (S.S. growth rate)	0.45	0.42	0.43	0.53	0.28	0.4
Shocks						
ρ_a (Productivity)	0.95	0.94	0.95	0.99	0.99	0.93
ρ_b (Risk premium)	0.19	0.31	0.39	0.2	0.15	0.19
ρ_g (Gov. spending)	0.97	0.98	0.98	0.98	0.95	0.98
ρ_i (Investment)	0.72	0.43	0.66	0.56	0.49	0.09
ρ_r (Monetary policy)	0.07	0.1	0.1	0.12	0.09	0.1
ρ_p (Price mark-up)	0.04	0.1	0.17	0.12	0.12	0.18
ρ_w (Wage mark-up)	0.97	0.33	0.38	0.87	0.18	0.1
ρ_{ga} (TFP impact on Gov.)	0.56	0.56	0.53	0.58	0.53	0.54
Shock St. Dev.						
η_a (TFP)	0.45	0.48	0.47	0.47	0.48	0.47
η_b (Risk premium)	2.14	2.87	3.16	2.57	3.3	3.6
η_g (Gov. spending)	0.57	0.5	0.5	0.51	0.5	0.51
η_i (Investment)	0.45	1.51	1.61	1.66	1.64	1.58
η_r (Monetary policy)	0.22	0.21	0.21	0.21	0.21	0.21
η_p (Price mark-up)	0.39	0.4	0.39	0.36	0.35	0.38
η_w (Wage mark-up)	0.18	0.56	0.57	0.47	0.56	0.58
$\eta_{\pi_{exp}}$ (Inflation expectations)	0.21	0.23	0.18	0.17	0.25	0.23
constant gain			0.044	0.005	0.03	0.006
(Log-) likl at mode	-1045.22	-977.92	-959.1	-981.96	-992.44	-990.68
MHM	-1156.11	-1057.68	-1032.82	-1074.43	-1072.53	-1067.11
Bayes' factor	0	42.74	53.54	35.47	36.3	38.65

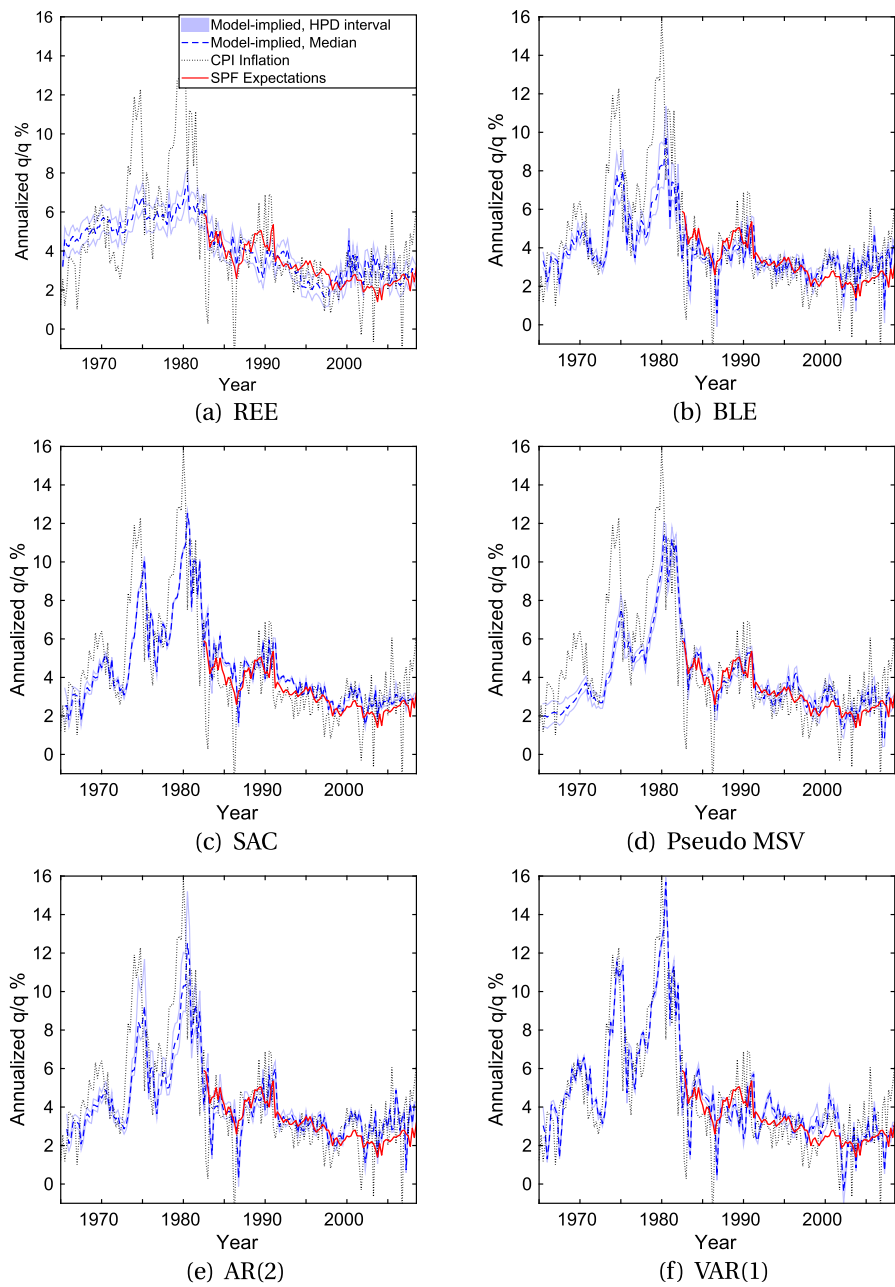


FIGURE 3. Model implied inflation expectations (blue), CPI inflation (black), and expectations from the SPF (red).

inflation/pre-Great Moderation era, without using any input on survey expectations over that period. This pattern is absent in the RE model, which is characterized by a more stable pattern for inflation expectations over the high inflation period. To distinguish how well each model tracks inflation survey expectations over the period where expecta-

TABLE 4. Correlations between survey- and model-generated inflation expectations and expectation errors. π_{t+1}^{SPF} denotes 1-quarter ahead inflation expectations from the SPF. $\mathbb{E}_t \pi_{t+1}$ denotes model-implied 1-quarter ahead inflation expectations. π_{t+1}^{data} denotes realized inflation at period $t + 1$.

Model	Correlation Between SPF and Model-Implied Inflation Expectations	Correlation Between Realized and Model-Implied Inflation Expectation Errors
	$\text{corr}(\pi_{t+1}^{\text{SPF}}, \mathbb{E}_t \pi_{t+1})$	$\text{corr}(\pi_{t+1}^{\text{data}} - \pi_{t+1}^{\text{SPF}}, \pi_{t+1}^{\text{data}} - \mathbb{E}_t \pi_{t+1})$
SAC	0.857	0.946
AR(2)	0.69	0.87
VAR(1)	0.371	0.798
BLE	0.496	0.837
REE	0.61	0.17
Pseudo MSV	0.59	0.818

tions data is available, we report two statistics for each model in Table 4. The first column reports the correlation between survey expectations π_{t+1}^{SPF} and model-implied inflation expectations $\mathbb{E}_t \pi_{t+1}$. SAC- and AR(2)-learning models yield the highest correlations and improve upon the REE benchmark, whereas the BLE, pseudo MSV-, and VAR(1)-learning models yield lower values compared to REE. Hence, in terms of capturing the *level* of inflation expectations, the REE benchmark is competitive and outperforms BLE and two of the learning models. The shortcoming of the REE model is its failure to capture expectation errors: in the second column of Table 4, we report the correlation between empirical inflation expectation errors $\pi_{t+1}^{\text{data}} - \pi_{t+1}^{\text{SPF}}$ (the difference between realized inflation and survey expectations) and model-implied expectation errors $\pi_{t+1}^{\text{data}} - \mathbb{E}_t \pi_{t+1}$ (the difference between realized inflation and model-implied inflation expectations). In this case, the REE benchmark yields a low correlation with 0.17, whereas BLE and learning models all yield higher values ranging between 0.8 and 0.95. Looking at both Tables 3 and 4 suggests that the SAC-learning model has the best fit in terms of inflation survey expectations.

To understand the dynamics around inflation expectations and distinguish the marginal contribution of learning dynamics, we plot the perceived mean and perceived persistence coefficients for the BLE and SAC-learning models in Figure 4. The equilibrium perception of inflation persistence β^* under the BLE model is 0.74. The time-varying perception in SAC-learning oscillates around the BLE-consistent value for most of the sample, starting to decline only after 2000 toward the end of the sample period. The main difference between BLE and SAC-learning comes from the perceived mean values; while the equilibrium value under BLE α^* is fixed at 0, the SAC-learning model displays a large degree of time variation in the mean. In particular, the high-inflation period of the 1970s and 1980s mainly transmits through the perceived mean in the learning model, which helps capture the inflation expectation dynamics better overall.

Our results are in line with Eusepi and Preston (2018a) and Eusepi, Giannoni, and Preston (2019), who show that beliefs under a constant-gain infinite-horizon learning

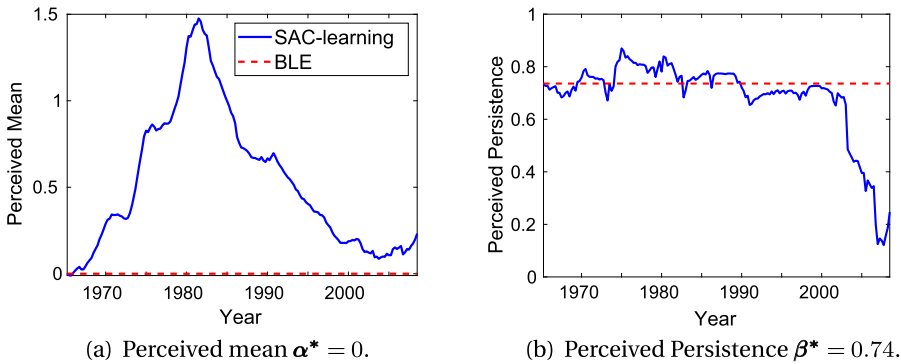


FIGURE 4. Belief coefficients α_t and β_t under SAC-learning, with BLE $\alpha^* = 0$ and $\beta^* = 0.74$.

approach fit U.S. data on inflation and interest rate expectations better than a rational expectations model. Our results confirm that learning dynamics continue to be important in capturing expectation dynamics when we replace the MSV-consistent PLM with an AR(1) heuristic.

Finally, we informally discuss the models' ability to capture movements in long-term inflation expectations, which generally remain firmly anchored in REE models even during periods of high and volatile inflation. Our BLE model suffers from the same shortcoming as REE models: since expectations are pinned down purely through the persistence coefficient β^* and the perceived mean is anchored at $\alpha^* = 0$, long-term inflation expectations remain stable in our BLE model. Given our median estimate of $\beta^* = 0.74$, expectations beyond 3 years remain firmly anchored regardless of the level of inflation. This is what distinguishes learning models from equilibrium models, where time-varying belief coefficients, in particular the perceived mean, can generate trend inflation and capture periods of deanchored long-term inflation expectations, as discussed in [Carvalho et al. \(2023\)](#).³⁴

4. POLICY APPLICATION: OPTIMAL SMOOTHING

In this section, we analyze the monetary policy implications for some of the estimated models.³⁵ A number of papers in the adaptive learning literature explore optimal monetary policy within the class of standard Taylor rule policies and look into the trade-off

³⁴Gaus and Gibbs (2018) suggest that Euler-equation learning models such as those considered in this paper produce better short-term inflation expectations. Infinite-horizon learning as in [Preston \(2005\)](#) and [Carvalho et al. \(2023\)](#) is more in line with long-run inflation expectations. They further note that infinite-horizon learning tends to improve the model fit more compared to Euler-equation learning. A more comprehensive horse race that includes infinite-horizon learning models is beyond the scope of our paper. Further note that [Carvalho et al. \(2023\)](#) report substantial improvements in fitting-inflation expectations data relative to baseline RE with their endogenous-gain learning model. We leave a comparison of this approach to BLE (and extensions thereof) to future work.

³⁵We leave the VAR(1)- and AR(2)-learning models out of this analysis and focus on the equilibrium models REE and BLE, against their learning counterparts SAC- and pseudo MSV-learning.

between inflation/output gap stabilization and central bank learning.³⁶ Our main focus in this section is the trade-off between interest rate smoothing and output/inflation stabilization, rather than the trade-off between inflation and output gap stabilization. Therefore, we fix the reaction coefficients on inflation, output gap and output gap difference at their estimated values and focus on the interest rate smoothing parameter ρ . Woodford (2003b) shows that under REE with forward-looking agents, optimal interest rate smoothing is typically high and close to unity across a wide range of specifications. In this section, we analyze how these results change with a backward-looking AR(1) rule under BLE and SAC-learning. Since our focus is on optimal interest rate smoothing, we use the following modified Taylor rule for monetary policy:

$$r_t = \rho r_{t-1} + \phi_\pi((1 - \rho)(\pi_t + \phi_y y_t) + \phi_{\Delta y} \Delta y_t) + \epsilon_t^r. \quad (4.1)$$

In the analysis below, we first fix the reaction parameters in all models at the estimated values under REE, $\phi_y = 0.11$, $\phi_{\Delta y} = 0.15$, and $\phi_\pi = 1.51$ in order to abstract away from any impact that the estimated parameter differences might have on the results. For the remaining parameters in BLE and REE, we leave the values at their posterior mean as reported in the baseline estimation Table 1. For the SAC- and pseudo MSV-learning cases, we use the parameter values associated with BLE and REE models, respectively, which helps us focus on disentangling the effects of learning from equilibrium models in isolation from the differences in the estimated parameter values. Furthermore, in order to prevent the presence of the projection facility in the learning models from affecting the optimal policy results, we fix the constant gain value in both models at a value of 0.001, which is sufficiently small to allow us to simulate the models without any projection facilities.³⁷

For this exercise, we use a grid of 500 points for the policy parameters ρ in each model, using a simulation length of 5000 periods in each case. For the BLE specification, we use $N = 200$ fixed-point iterations to calculate the equilibrium values β^* for each value of the policy parameter, as in the likelihood evaluation in Section 3.2. The number of periods is sufficient to ensure convergence of the learning parameters. In order to avoid any effects of the transient learning dynamics, we discard the initial 80% of the sample in each simulation and use the remaining 20% (1000 periods) to compute the associated moments of inflation, output gap, and interest rate.

Figure 5 reports the percentage change in the standard deviations of the output gap, inflation, and interest rate as a function of the interest rate smoothing parameter ρ . Under REE and pseudo MSV-learning models, smoothing is beneficial in terms of stabilizing variation in the output gap, y_t , and inflation, π_t , up to a point. Under BLE and SAC-learning specifications, we observe a different pattern where the stabilizing effects

³⁶A nonexhaustive list includes Orphanides and Williams (2005, 2006, 2008), Evans and Honkapohja (2003), Preston (2006), and Gasteiger (2014).

³⁷Different gain values can also have important implications on the optimal parameters in learning models, as shown in Orphanides and Williams (2004). Our main focus in this section is how the degree of knowledge about the underlying system under BLE affects the monetary policy implications relative to REE. Therefore, we abstract away from such considerations.

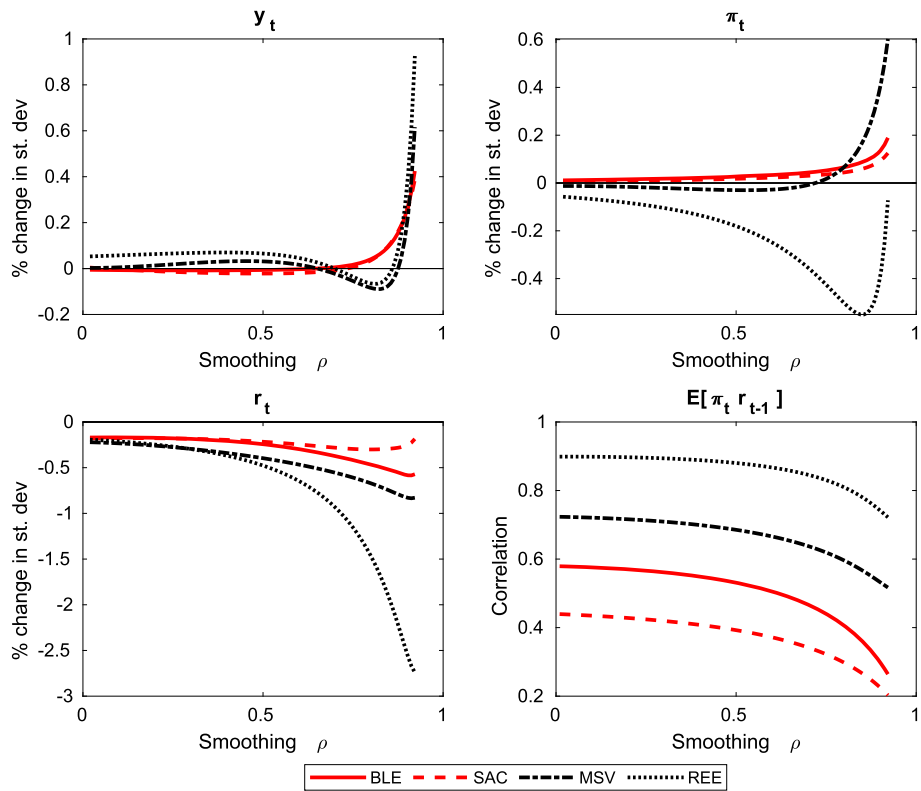


FIGURE 5. Standard deviations and correlation between inflation and lagged interest rate (y-axis) as a function of interest rate smoothing ρ (x-axis).

disappear and both inflation and output gap become more volatile as the smoothing parameter increases. To formalize this, we introduce an ad hoc loss-function $E[L]$ in terms of the discounted sum of weighted squared inflation, output gap growth, and interest rate:

$$E[L] = (1 - \vartheta)E\left[\sum_{t=0}^{\infty} \vartheta^t [\omega_{\pi}\pi_t^2 + \omega_y\Delta y_t^2 + \omega_r r_t^2]\right] = \omega_{\pi}\sigma_{\pi}^2 + \omega_y\sigma_{\Delta y}^2 + \omega_r\sigma_r^2, \tag{4.2}$$

with ω_{π} , ω_y , and ω_r the weights on inflation, the growth of the output gap, and the interest rate, respectively. In this paper, following the approach in [Slobodyan and Wouters \(2012a\)](#), we model the output gap as the deviation of output \tilde{y}_t from the underlying productivity process ϵ_t^a , that is, $y_t = \tilde{y}_t - \Phi_p \epsilon_t^a$ with Φ_p the estimated value of production of fixed costs for each model.

Table 5 reports the optimal smoothing values ρ^* for 3 combinations of these weights, where we normalize $\omega_{\pi} = 1$. The optimal smoothing ρ^* under BLE and SAC-learning is lower than REE and pseudo MSV-learning models for all combinations, and the REE model always yields the highest optimal ρ^* . Of particular interest is the point where the weight on nominal interest rate stabilization in the objective function is zero, $\omega_r = 0$. In this case, the BLE model implies an optimal smoothing equal to 0.

TABLE 5. Optimal smoothing parameter for some cases.

Model	ω_π	ω_y	ω_r	Optimal ρ^*
REE	1	0.048	0	0.91
	1	0.048	0.1	0.92
	1	0.1	0.1	0.91
BLE	1	0.048	0	0
	1	0.048	0.1	0.79
	1	0.1	0.1	0.82
SAC	1	0.048	0	0.6
	1	0.048	0.1	0.79
	1	0.1	0.1	0.77
pseudo MSV	1	0.048	0	0.75
	1	0.048	0.1	0.82
	1	0.1	0.1	0.84

One reason for this result is that backward-looking agents do not consider the movements in the interest rate when forming their expectations. As the smoothing coefficient increases, the contemporaneous reaction of the interest rate to inflation and the output gap decreases. Agents do not internalize future movements of the interest rate.³⁸ As a result, higher smoothing is interpreted as a weaker reaction to inflation and output-growth fluctuations on their part, which leads to higher volatility in inflation and output gap. Since agents do not internalize the stabilizing effect of the policy rate smoothing (as would be the case under REE and pseudo MSV-learning), fluctuations in the policy rate **may** become less and less costly, thereby resulting in less smoothing. A similar argument applies to the SAC learning model. But the main reason behind the substantially lower smoothing under the BLE and SAC-learning lies in the persistence inherent in the model when agents are purely backward looking. To see that, consider the 3-equation purely forward-looking NK model in (2.30) with the following simple Taylor rule, where for the sake of exposition we assume that the central bank targets inflation only:

$$i_t = \rho i_{t-1} + \phi_\pi \pi_t. \quad (4.3)$$

Considering the REE model, by iterating the above interest rate rule backwards and using the forward-looking Phillips curve, the rule writes as follows:

$$i_t = \frac{\phi_\pi \gamma}{1 - \rho \lambda} \sum_{s=0}^{\infty} \lambda^s y_{t+s+1} + \frac{\phi_\pi \gamma \rho}{1 - \rho \lambda} \sum_{s=0}^{\infty} \rho^{s-1} y_{t-s}. \quad (4.4)$$

³⁸Under rational expectations, agents make forecasts of interest rates into the infinite future. In the BLE and related learning models, interest rate forecasts play no role. In this adaptive learning model, there is no transmission mechanism through the term structure of interest rates, only current interest rates matter. Extending BLE adaptive learning models with a transmission mechanism through the term structure is an interesting avenue for future research.

As argued by [Giannoni \(2014\)](#), the optimal monetary policy under commitment in a purely forward-looking model results in a bounded solution where the endogenous variables depend not only upon expected future values of disturbances, but also on pre-determined variables. This means that optimal policy introduces history dependence, something that is missing in simple interest rate rules without smoothing and pure inflation targeting. More importantly, [Giannoni \(2014\)](#) shows that an optimal interest rate rule that is not only inertial but also superinertial can be derived from the first-order conditions of the optimal policy problem of the central bank. As (4.4) reveals, interest rate smoothing, captured by ρ , is necessary in order to introduce history dependence in a purely forward-looking model. Clearly, setting $\rho = 0$ in (4.4) shuts down dependence on past data and makes the rule implicitly purely forward-looking in nature. This is why the REE requires a higher smoothing parameter.

Let us now consider BLE or SAC learning in the same simple 3-equation NK model with the above rule (4.3) but now without smoothing (i.e., $\rho = 0$). In this case, the Phillips curve after plugging inflation expectations (assuming zero mean in inflation expectations) takes the following form:

$$\pi_t = \lambda\beta\pi_{t-1} + \gamma y_t. \quad (4.5)$$

Plugging the above expression in (4.4) and iterating backwards, we get

$$i_t = \phi_\pi \gamma \sum_{s=0}^{\infty} (\lambda\beta)^s y_{t-s}. \quad (4.6)$$

As equation (4.6) reveals, the backward-looking nature of expectations introduces persistence in the model that makes the interest rate depend on current and past information only. As such, interest rate smoothing is not necessary, nor does it add further information in interest rate setting. That explains why our simulations find that zero or substantially lower smoothing is required under the BLE or SAC learning.

In the literature, the observed rate of interest rate smoothing in the historical data has been attributed to the presence of forward-looking agents ([Woodford \(2003b\)](#)), where a high degree of smoothing helps introduce history dependence into agents' beliefs and steers private-sector expectations of future policy in the right direction. High interest rate smoothing or first difference rules have also been found beneficial in models with central bank uncertainty and learning about the data or model parameters ([Sack and Wieland \(2000\)](#)), as well as in studies where both agents and the central bank use adaptive learning ([Orphanides and Williams \(2007\)](#), [Woodford \(2013\)](#)).³⁹ Our results here suggest that smoothing is not desirable with boundedly rational agents that use backward-looking forecasting rules in the absence of central bank learning. Central bank learning could affect the resulting optimal interest rate inertia in either direction, in the presence of backward-looking learning rules adopted by the private sector of the economy. We leave a further exploration of this topic to future research.

³⁹[Eusepi and Preston \(2018b\)](#) provide an extensive and very detailed review on the properties of interest rate rules under imperfect knowledge.

5. CONCLUDING REMARKS

In this paper, we generalize the BLE concept with optimal AR(1) beliefs to an n -dimensional linear stochastic framework and provide an approximation and estimation method for it. We apply the concept to a simple NK model to derive analytical results and build intuition. We then estimate BLE in the workhorse Smets and Wouters (2007) model and compare the in-sample fit and out-of-sample forecasting performance of different learning models. In this way, we disentangle the effects of the degree of knowledge about the underlying economy and of learning on the model fit. We find that replacing the cross-restrictions of REE with those implied by BLE plays an important role in improving in-sample fitness and pseudo out-of-sample forecasting performance up to 4 quarters. Introducing learning with AR(1) expectations improves the fitness further, particularly when the model is reestimated with short-term inflation expectations from survey data. In particular, SAC-learning with AR(1) beliefs provides the best fit among the constant-gain learning models considered in this paper when short-term survey data on inflation expectations are taken into account.

Our work opens up several important avenues of future research. First, our results call attention to the general class of restricted perceptions equilibria that consider different degrees of misspecification and accompanying solution algorithms to empirically estimate these equilibria. Second, sample-autocorrelation learning, which is based on a method-of-moments estimator for the AR(1) rule, should be extended and generalized as an alternative to the constant-gain recursive least squares learning in order to account for any class of PLM and to complement the corresponding restricted perceptions equilibrium concepts. In general, estimation methods of optimal forecasting heuristics within macroeconomic models seem a plausible and empirically relevant avenue for future work. Policy analysis under optimal forecasting heuristics is an important application of these theoretical and empirical tools. Finally, while the empirical horse race in this paper is limited to Euler-equation learning models, extending the analysis to other approaches such as infinite-horizon learning is an important topic for future work.

REFERENCES

- Adam, Klaus (2007), “Experimental evidence on the persistence of output and inflation.” *The Economic Journal*, 117 (520), 603–636. [1402, 1406, 1409, 1416]
- Adam, Klaus and Mario Padula (2011), “Inflation dynamics and subjective expectations in the United States.” *Economic Inquiry*, 49 (1), 13–25. [1404]
- Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry (2021), “Imperfect macroeconomic expectations: Evidence and theory.” *NBER Macroeconomics Annual*, 35 (1), 1–86. [1421]
- Angeletos, George-Marios and Chen Lian (2016), “Incomplete information in macroeconomics: Accommodating frictions in coordination.” In *Handbook of Macroeconomics*, Vol. 2, 1065–1240, Elsevier, Amsterdam: North-Holland. [1404]

Assenza, Tiziana, Peter Heemeijer, Cars H. Hommes, and Domenico Massaro (2021), “Managing self-organization of expectations through monetary policy: A macro experiment.” *Journal of Monetary Economics*, 117, 170–186. [1406, 1409, 1416]

Audzei, Volha and Sergey Slobodyan (2022), “Sparse restricted perceptions equilibrium.” *Journal of Economic Dynamics and Control*, 139, 104415. [1402, 1406]

Berardi, Michele and Jaqueson K. Galimberti (2017), “Empirical calibration of adaptive learning.” *Journal of Economic Behavior & Organization*, 144, 219–237. [1405, 1423]

Beshears, John, James J. Choi, Andreas Fuster, David Laibson, and Brigitte C. Madrian (2013), “What goes up must come down? Experimental evidence on intuitive forecasting.” *American Economic Review*, 103 (3), 570–574. [1406]

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (2020), “Overreaction in macroeconomic expectations.” *American Economic Review*, 110 (9), 2748–2782. [1421]

Branch, William A. (2004), “Restricted perceptions equilibria and learning in macroeconomics.” In *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*, 135–160. [1402, 1409]

Branch, William A. and George W. Evans (2006), “A simple recursive forecasting model.” *Economics Letters*, 91 (2), 158–166. [1404]

Branch, William A. and George W. Evans (2010), “Asset return dynamics and learning.” *The review of financial studies*, 23 (4), 1651–1680. [1402]

Branch, William A., George W. Evans, and Bruce McGough (2014), “Perpetual learning and stability in macroeconomic models.” Report, University of Oregon. [1421]

Bray, Margaret (1982), “Learning, estimation, and the stability of rational expectations.” *Journal of Economic Theory*, 26 (2), 318–339. [1419]

Brock, William A. and Anastasios G. Malliaris (1989), *Differential Equations, Stability and Chaos in Dynamic Economics*. Advanced Textbooks in Economics (90A16 BROd). [1414]

Bullard, James (2006), “The learnability criterion and monetary policy.” *Review-Federal Reserve Bank of Saint Louis*, 88 (3), 203. [1404]

Bullard, James, George W. Evans, and Seppo Honkapohja (2008), “Monetary policy, judgment, and near-rational exuberance.” *American Economic Review*, 98 (3), 1163–1177. [1402]

Bullard, James and Kaushik Mitra (2002), “Learning about monetary policy rules.” *Journal of Monetary Economics*, 49 (6), 1105–1129. [1406, 1415, 1416]

Canova, Fabio and Luca Gambetti (2010), “Do expectations matter? The great moderation revisited.” *American Economic Journal: Macroeconomics*, 2 (3), 183–205. [1404]

Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston (2023), “Anchored inflation expectations.” *American Economic Journal: Macroeconomics*, 15 (1), 1–47. [1424, 1430, 1434]

Chung, Hyein and Wei Xiao (2013), “Cognitive consistency, signal extraction, and macroeconomic persistence.” Report, SUNY Binghamton. [1408]

Clarida, Richard, Jordi Gali, and Mark Gertler (1999), “The science of monetary policy: A New Keynesian perspective.” *Journal of Economic Literature*, 37 (4), 1661–1707. [1417]

Clark, Todd E. and Kenneth D. West (2007), “Approximately normal tests for equal predictive accuracy in nested models.” *Journal of Econometrics*, 138 (1), 291–311. [1407]

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar (2018), “The formation of expectations, inflation, and the Phillips curve.” *Journal of Economic Literature*, 56 (4), 1447–1491. [1402]

Croushore, Dean (1993), “Introducing: the survey of professional forecasters.” *Business Review-Federal Reserve Bank of Philadelphia*, 6, 3. [1405, 1424]

DeCanio, Stephen J. (1979), “Rational expectations and learning from experience.” *The Quarterly Journal of Economics*, 93 (1), 47–57. [1419]

Del Negro, Marco and Stefano Eusepi (2011), “Fitting observed inflation expectations.” *Journal of Economic Dynamics and Control*, 35 (12), 2105–2131. [1405]

Elaydi, Saber (2005), *An Introduction to Difference Equations*, third edition. Springer, New York. [1413]

Enders, Walter (2008), *Applied Econometric Time Series*. John Wiley and Sons, New York. [1407]

Eusepi, Stefano, Marc Giannoni, and Bruce Preston (2019), “On the limits of monetary policy.” Unpublished paper, University of Texas, Austin. [1433]

Eusepi, Stefano and Bruce Preston (2018a), “Fiscal foundations of inflation: Imperfect knowledge.” *American Economic Review*, 108 (9), 2551–2589. [1433]

Eusepi, Stefano and Bruce Preston (2018b), “The science of monetary policy: An imperfect knowledge perspective.” *Journal of Economic Literature*, 56 (1), 3–59. [1438]

Evans, George (1985), “Expectational stability and the multiple equilibria problem in linear rational expectations models.” *The Quarterly Journal of Economics*, 100 (4), 1217–1233. [1419]

Evans, George W. and Seppo Honkapohja (2001), *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ. [1402, 1404, 1405, 1414]

Evans, George W. and Seppo Honkapohja (2003), “Expectations and the stability problem for optimal monetary policies.” *The Review of Economic Studies*, 70 (4), 807–824. [1405, 1406, 1435]

Evans, George W. and Seppo Honkapohja (2006), “Monetary policy, expectations and commitment.” *Scandinavian Journal of Economics*, 108 (1), 15–38. doi:10.1111/j.1467-9442.2006. [1406]

Fuhrer, Jeff (2017), “Expectations as a source of macroeconomic persistence: Evidence from survey expectations in a dynamic macro model.” *Journal of Monetary Economics*, 86, 22–35. [1405]

Fuhrer, Jeffrey C. (2006), “Intrinsic and inherited inflation persistence.” Working Paper No. 05-8, Federal Reserve Bank Boston. [1417]

Fuhrer, Jeffrey C. (2010), “Inflation persistence.” In *Handbook of Monetary Economics*, Vol. 3, 423–486, Elsevier, Amsterdam: North-Holland. [1417]

Fuster, Andreas, Benjamin Hebert, and David Laibson (2010), “Investment dynamics with natural expectations.” *International Journal of Central Banking*, 8 (81), 243. [1406, 1421]

Fuster, Andreas, Benjamin Hebert, and David Laibson (2012), “Natural expectations, macroeconomic dynamics, and asset pricing.” *NBER Macroeconomics Annual*, 26 (1), 1–48. [1406]

Fuster, Andreas, David Laibson, and Brock Mendel (2010), “Natural expectations and macroeconomic fluctuations.” *Journal of Economic Perspectives*, 24 (4), 67–84. [1402, 1406]

Gali, Jordi (2008), *Introduction to Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, Princeton, NJ. [1414]

Gaspar, Vitor, Frank Smets, and David Vestin (2010), “Inflation expectations, adaptive learning and optimal monetary policy.” In *Handbook of Monetary Economics*, Vol. 3, 1055–1095, Elsevier, Amsterdam: North-Holland. Chapter 19. [1407]

Gasteiger, Emanuel (2014), “Heterogeneous expectations, optimal monetary policy, and the merit of policy inertia.” *Journal of Money, Credit and Banking*, 46 (7), 1535–1554. [1435]

Gaus, Eric and Christopher G. Gibbs (2018), “Expectations and the empirical fit of DSGE models.” Report. [1402, 1405, 1419, 1423, 1428, 1434]

Gelain, Paolo, Nikolay Iskrev, Kevin J. Lansing, and Caterina Mendicino (2019), “Inflation dynamics and adaptive expectations in an estimated DSGE model.” *Journal of Macroeconomics*, 59, 258–277. [1406]

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer (2016), “Expectations and investment.” *NBER Macroeconomics Annual*, 30 (1), 379–431. [1405]

Giannoni, Marc P. (2014), “Optimal interest-rate rules and inflation stabilization versus price-level stabilization.” *Journal of Economic Dynamics and Control*, 41, 110–129. [1404, 1438]

Greenberg, Edward (2012), *Introduction to Bayesian Econometrics*. Cambridge University Press. [1425]

Hamilton, James (1994), *Time Series Econometrics*. Princeton University Press, Princeton, NJ. [1409]

Hommes, Cars (2021), “Behavioral and experimental macroeconomics and policy analysis: A complex systems approach.” *Journal of Economic Literature*, 59 (1), 149–219. [1409]

Hommes, Cars and Joep Lustenhouwer (2019), “Inflation targeting and liquidity traps under endogenous credibility.” *Journal of Monetary Economics*, 107, 48–62. [1406]

Hommes, Cars, Kostas Mavromatis, Tolga Özden, and Mei Zhu (2023), “Supplement to ‘Behavioral learning equilibria in New Keynesian models.’” *Quantitative Economics Supplemental Material*, 14, <https://doi.org/10.3982/QE1533>. [1412]

Hommes, Cars and Gerhard Sorger (1998), “Consistent expectations equilibria.” *Macroeconomic Dynamics*, 2 (3), 287–321. [1402, 1410]

Hommes, Cars and Mei Zhu (2014), “Behavioral learning equilibria.” *Journal of Economic Theory*, 150, 778–814. [1402, 1407, 1408, 1409, 1410, 1411, 1413]

Huang, Kevin X. D., Zheng Liu, and Tao Zha (2009), “Learning, adaptive expectations and technology shocks.” *Economic Journal*, 119 (536), 377–405. [1404]

Jääskelä, Jarkko and Rebecca McKibbin (2010), “Learning in an estimated small open economy model.” Research Discussion Paper (2010-02), Reserve Bank of Australia. [1422, 1428]

Lansing, Kevin J. (2009), “Time-varying US inflation dynamics and the New Keynesian Phillips curve.” *Review of Economic Dynamics*, 12 (2), 304–326. [1402, 1406]

Lansing, Kevin J. and Jun Ma (2017), “Explaining exchange rate anomalies in a model with Taylor-rule fundamentals and consistent expectations.” *Journal of International Money and Finance*, 70, 62–87. [1402, 1406]

Lucas, Robert E. (1972), “Expectations and the neutrality of money.” *Journal of Economic Theory*, 4 (2), 103–124. [1402]

Marcet, Albert and Juan P. Nicolini (2003), “Recurrent hyperinflations and learning.” *American Economic Review*, 93 (5), 1476–1498. [1404]

Marcet, Albert and Thomas J. Sargent (1989), “Convergence of least squares learning mechanisms in self-referential linear stochastic models.” *Journal of Economic Theory*, 48 (2), 337–368. [1405]

Milani, Fabio (2005), “Adaptive learning and inflation persistence.” University of California, Irvine—Department of Economics. [1422, 1424]

Milani, Fabio (2007), “Expectations, learning and macroeconomic persistence.” *Journal of Monetary Economics*, 54 (7), 2065–2082. [1405, 1425]

Milani, Fabio (2011), “Expectation shocks and learning as drivers of the business cycle.” *The Economic Journal*, 121 (552), 379–401. [1405]

Muth, John F. (1961), “Rational expectations and the theory of price movements.” *Econometrica: Journal of the Econometric Society*, 315–335. [1402]

Nelson, Charles R. (1972), “The prediction performance of the FRB-MIT-PENN model of the US economy.” *The American Economic Review*, 62 (5), 902–917. [1407]

Ormeño, Arturo and Krisztina Molnár (2015), “Using survey data of inflation expectations in the estimation of learning and rational expectations models.” *Journal of Money, Credit and Banking*, 47 (4), 673–699. [1406, 1430]

Orphanides, Athanasios and John Williams (2004), “Imperfect knowledge, inflation expectations, and monetary policy.” In *The Inflation-Targeting Debate*, 201–246, University of Chicago Press, Chicago, IL. [1405, 1435]

Orphanides, Athanasios and John C. Williams (2005), “Inflation scares and forecast-based monetary policy.” *Review of Economic Dynamics*, 8 (2), 498–527. [1406, 1435]

Orphanides, Athanasios and John C. Williams (2006), “Monetary policy with imperfect knowledge.” *Journal of the European Economic Association*, 4 (2–3), 366–375. [1435]

Orphanides, Athanasios and John C. Williams (2007), “Robust monetary policy with imperfect knowledge.” *Journal of monetary Economics*, 54 (5), 1406–1435. [1438]

Orphanides, Athanasios and John C. Williams (2008), “Learning, expectations formation, and the pitfalls of optimal control monetary policy.” *Journal of Monetary Economics*, 55, S80–S96. [1435]

Pfajfar, Damjan and Blaž Žakelj (2014), “Experimental evidence on inflation expectation formation.” *Journal of Economic Dynamics and Control*, 44, 147–168. [1409, 1416]

Preston, Bruce (2005), “Learning about monetary policy rules when long-horizon expectations matter.” *International Journal of Central Banking*. [1405, 1419, 1434]

Preston, Bruce (2006), “Adaptive learning, forecast-based instrument rules and monetary policy.” *Journal of Monetary Economics*, 53 (3), 507–535. [1406, 1435]

Sack, Brian and Volker Wieland (2000), “Interest-rate smoothing and optimal monetary policy: A review of recent empirical evidence.” *Journal of Economics and Business*, 52 (1–2), 205–228. [1438]

Sargent, Thomas, Noah Williams, and Tao Zha (2009), “The conquest of South American inflation.” *Journal of Political Economy*, 117 (2), 211–256. [1404]

Sargent, Thomas J. (1991), “Equilibrium with signal extraction from endogenous variables.” *Journal of Economic Dynamics and Control*, 15 (2), 245–273. [1402]

Sims, Christopher (1999), “Matlab optimization software.” QM&RBC Codes. [1424]

Slobodyan, Sergey and Raf Wouters (2012a), “Learning in a medium-scale DSGE model with expectations based on small forecasting models.” *American Economic Journal: Macroeconomics*, 4 (2), 65–101. [1406, 1420, 1422, 1423, 1428, 1430, 1436]

Slobodyan, Sergey and Raf Wouters (2012b), “Learning in an estimated medium-scale DSGE model.” *Journal of Economic Dynamics and Control*, 36 (1), 26–46. [1422, 1423, 1424, 1425, 1428, 1430]

Smets, Frank and Rafael Wouters (2007), “Shocks and frictions in US business cycles: A Bayesian DSGE approach.” *American Economic Review*, 97 (3), 586–606. [1401, 1402, 1403, 1406, 1419, 1424, 1428, 1439]

Stock, James H. and Mark W. Watson (2007), “Why has US inflation become harder to forecast?” *Journal of Money, Credit and Banking*, 39, 3–33. [1407]

Williams, Noah (2003), “Adaptive learning and business cycles.” Report, Princeton University. [1404]

Woodford, Michael (2003a), *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ. [1403, 1414, 1417]

Woodford, Michael (2003b), “Optimal interest-rate smoothing.” *The Review of Economic Studies*, 70 (4), 861–886. [1435, 1438]

Woodford, Michael (2013), “Macroeconomic analysis without the rational expectations hypothesis.” *Annu. Rev. Econ.*, 5 (1), 303–346. [1404, 1438]

Xiao, Wei and Junyi Xu (2014), “Expectations and optimal monetary policy: A stability problem revisited.” *Economics Letters*, 124 (2), 296–299. [1421]

Co-editor Tao Zha handled this manuscript.

Manuscript received 24 January, 2020; final version accepted 26 April, 2023; available online 16 May, 2023.