A Dynare Toolbox for Social Learning Expectations*

Alex Grimaud† Isabelle Salle‡ Gauthier Vermandel§

January 8, 2024

Abstract

Social learning (SL) is a behavioral model in which expectations and the resulting aggregate dynamics stem from the interactions of a large number of heterogeneous agents. Nonetheless, this framework has so far lacked a parsimonious development with a general-solution method. This paper bridges this gap and introduces a Dynare toolbox to solve any linear state-space model with SL expectations, opening up a wide range of potential applications. As an illustration, optimal monetary policy rules are studied in a microfounded New Keynesian (NK) model under SL and rational expectations (RE).
***Keywords*:** Inflation targeting, Monetary policy and Heterogeneous expectations.
***JEL codes*:** E32, E52, E58 and E71.

1 Introduction

A major contribution of the late Prof. Jasmina Arifovic is the application of genetic algorithms (GAs), initially developed for optimization purposes, to model heterogeneous decisions and expectations in economic models.[1] One fruitful application is the use of these algorithms to model the evolution of heterogeneous beliefs under social learning (SL). The observed heterogeneity of real-world expectations[2] makes this topic a particularly relevant one, not least because of the key role such expectations play in the transmission channel of

*The present work has benefited from the financial support of the Austrian National Bank (OeNB) grant No. 18611.

†Department of Economics, Vienna University of Economics and Business, Vienna, AT and Institute of Statistics and Mathematical Methods in Economics, TU Wien, Vienna, AT (alex.grimaud@wu.ac.at).

‡Department of Economics, University of Ottawa, Ontario, Canada; Amsterdam School of Economics, University of Amsterdam and Tinbergen Institute, NL (isalle@uottawa.ca).

§École Polytechnique UMR 7641, Paris, FR and Université Paris-Dauphine UMR 8007, Paris, FR (gauthier@vermandel.fr).

[1] This class of learning models is based on the evolutionary principle of 'survival of the fittest' among large populations [Holland, 1975]. See Arifovic [2000] for a survey of early contributions and Arifovic et al. [2013, 2018, forthcoming] for recent examples.

[2] This is true both in lab and survey data, and no matter whether expectations are elicited from professional forecasters, households, firms, managers or policymakers themselves; see, *inter alia*, Coibion and Gorodnichenko [2015], Hommes [2021], Weber et al. [2022]. Meeks and Monti [2023] find that the cross-section dispersion in inflation expectations plays a non-negligible role in inflation dynamics during the 2007-2009 crisis.

macroeconomic policies. Disagreement between agents may arise for a variety of reasons: for instance, if they have different cognitive abilities, rely on different information sets or use different forecasting models.

However, integrating SL expectations into standard macroeconomic models is challenging because they introduce a non-linear and backward-looking learning process into otherwise standard log-linearized dynamic stochastic general equilibrium (DSGE) models. Throughout Jasmina's career, this difficulty was addressed by agent-level simulations of backward-looking systems. This is the so-called temporary equilibrium approach, as employed in many other learning applications in macroeconomics.[3] This implementation, however, makes it hard to scale up the models under SL expectations or to bring them to the data with full-information methods, thereby limiting the reach of her work within mainstream macroeconomic circles. On a technical note, the temporary equilibrium approach implies that the policy function is not time-invariant, resulting in a violation of the law of iterated projections, which is itself an essential element of the microfoundations of forward-looking models. This inconsistency was relatively innocuous in the context of the early learning literature, which focused on equilibrium selection and asymptotic behaviors of learning economies.[4] However, it could represent a more serious hurdle in recent works aiming to explain the business cycle with perpetual learning dynamics.

Anecdotally, Jasmina had been invited several times to the Bank de France as a visiting scholar by Michel Julliard. Michel is a core member of the team behind `Dynare`, the most commonly used software platform to work with standard DSGE models [Juillard et al., 1996]. The two spoke multiple times about introducing SL mechanisms into this software to address the aforementioned limitations of the existing SL models, but the project never took off. In this special issue, we pay tribute to our friend and co-author by completing this project.

In order to achieve this, we reformulate SL expectations to obtain a parsimonious learning model using a general solution method. We are then able to apply the tools of the benchmark rational expectations (RE) framework to any DSGE model under SL expectations. The result is a (micro-founded) heterogeneous-expectation New Keyne-

[3]See, in particular, the adaptive learning literature [Evans and Honkapohja, 2001]. Heuristic-switching models are another example of this temporary equilibrium approach [Hommes, 2021].

[4]See, *inter alia*, Evans [1985], Marcet and Sargent [1989], Evans and Honkapohja [1995], Arifovic et al. [2013]. This literature focuses on whether agents may learn to coordinate on rational expectation equilibria (REE) in environments where there may be a multiplicity of them. Whenever it is the case, the issue of the time-invariance of the policy function is resolved by construction.

sian (HENK) framework[5] deployed in a `Dynare` toolbox.[6] Our reformulation and the associated toolbox make it possible to introduce heterogeneous SL expectations into any economic or finance model, no matter its scale, and to estimate the resulting model using Bayesian methods for counterfactual analysis. In this paper, we provide general functions to simulate time series and impulse response functions (IRFs). The toolbox also allows for welfare analysis within the context of the baseline NK model; we provide an example of how this can be done.

We make two key assumptions to develop our framework. First, we assume that SL agents understand and agree not only about the structure of the economy, but also about the effects of shocks and lagged endogenous variables on the variables that they have to forecast. However, they are uncertain about the long-run values or 'long-term drifts' of these variables.[7] They hold heterogeneous beliefs about these drifts, and these beliefs evolve according to the two operators underlying SL, namely i) the individual exploration of the space of strategies (the mutation operator) and ii) the collective exploitation of the existing strategies (the selection operator via a tournament). An economic interpretation of these mechanisms could be a setting in which agents receive idiosyncratic news (or private signals) and update their forecasts by copying the most accurate ones among their peers. Such copying would take place as part of social interactions in which pairs of individuals exchange ideas and opinions. When referring to the accuracy of the resulting forecasts, we mean the accuracy of the predictions derived from a PLM using a particular intercept value over the (recent) history.

The second important assumption that we use is that agents understand that learning influences the economy in the current and future periods. We refer to this assumption as 'internal rationality' because agents internalize the effects of the deviations from RE that result from their learning behaviors in their forecasting model of the economy [Adam and Marcet, 2011]. This means that agents form expectations about the future deviations of the aggregate expectations from RE, i.e. about the future biases of aggregate expectations with respect to the RE benchmark, in order to form their own expectations of the endogenous variables.

[5]HENK models are different from the TANK or HANK (for 'two-agent' or 'heterogeneous-agent') models. In HENK models, heterogeneity concerns mostly expectations, while in HANK models, it usually regards idiosyncratic variables such as wealth, risk, or productivity [see e.g Kaplan et al., 2018, Bilbiie, 2020, Ottonello and Winberry, 2020]. Moreover, in HENK models, beliefs' heterogeneity might be dynamic and the distribution of the different beliefs may vary over time [see e.g Branch and McGough, 2010, Massaro, 2013, Andrade et al., 2019]. Finally, in HANK models, it is interesting to remark that due to the complexity of their policy functions, beliefs about future variables are often assumed to be partially deterministic or built under certainty equivalence.

[6]Here, we will refer to the citation of the toolbox and the associated DOI.

[7]We borrow this expression to Eusepi and Preston [2011, 2018b], Carvalho et al. [2023], who use a similar assumption, albeit in a representative-agent framework. It differs from the 'steady-state learning' implemented in non-linear models where agents forecast using a mis-specified (under-parametrized) rule that only involves an intercept; see, e.g., Evans et al. [2008]. We use a log-linearized model and the PLM of the agents is correctly specified.

Detailing our approach by using the vocabulary of the learning literature, SL agents learn about the intercept of the minimum state variable (MSV) solution of the model. This MSV solution also includes an additional state variable representing the deviations of aggregate expectations from RE. This implies that both the agents' perceived laws of motion (PLMs) and the actual law of motion (ALM) of the economy internalize the effects of learning on endogenous variables and, hence, that the PLMs are well-specified. Therefore, in our setting, SL expectations result in '*biased rational forecasts with internal rationality*,' where individual steady-state beliefs are biased, heterogeneous and evolve under SL. The bias of aggregate expectations becomes an additional state variable of the model which may differ from zero, which would be its value in the absence of learning and under RE. We then define the information sets that the SL agents and an RE observer use to forecast these future aggregate biased beliefs in the economy. Using only these minimal assumptions, this paper then shows how to write any SL model in the standard, reduced form common in the DSGE literature, and how to implement it in `Dynare`.

Restricting SL to the long-term drifts of the forward-looking variables offers several advantages.[8] First, it allows us to implement SL expectations in larger models than would otherwise be possible. This is because having agents simultaneously learn all the coefficients of the MSV solution quickly becomes computationally prohibitive and too complex when more state variables are introduced. Second, this formulation nests the RE benchmark, which eases model comparisons. Third, in terms of economic interpretation, this formulation allows us to model particularly intuitive phenomena. One such case is unanchored inflation expectations, such as can result from imperfect central bank (CB) credibility or a lack of public knowledge of the target. Another case is that of animal spirits, which for our purposes could be interpreted as SL output gap expectations becoming positively or negatively biased on average (reflecting optimism and pessimism, respectively). A further possibility concerns disagreement about the fundamental value of asset prices in a model featuring them. In all of these cases, the relevant dynamics are unable to be simulated under RE because all of the variables mean-revert towards their REE values, which are zero by construction. As for the assumption of internal rationality, it allows us to depart from the aforementioned temporary equilibrium solution and establish a time-invariant policy function that ensures that the law of iterated expectations holds in aggregate under SL.

The rest of the paper is organized as follows. Section 2 spells out our reformulation of SL within the three-equation NK environment used in Arifovic, Bullard, and Kostyshyna [2013] and discusses the main assumptions and implications of our approach. Section 3 generalizes the approach to any model with forward-looking variables. This section is

[8] Eusepi and Preston [2018a] further motivate this approach by emphasizing that it accounts for properties of survey forecast data [Kozicki and Tinsley, 2012] and generates the quantitatively relevant expectation-feedback dynamics in learning models.

useful for providing an overview of the Dynare toolbox. Section 4 provides an application of this toolbox to optimal monetary policy under SL versus RE in the three-equation NK model with history-dependent interest-rate rules. We show how heterogeneous expectations create dispersion in prices and consumption levels. Such dispersion entails an additional welfare loss with respect to the RE benchmark with exogenous fluctuations only, therefore complicating the stabilization trade-off faced by the CB. We find that a history-dependent monetary policy rule with an aggressive coefficient on inflation developments is optimal under SL. Moreover, such a policy rule is robust to the type of expectations actually employed by agents; no matter whether policies are optimized assuming RE or SL, they are welfare-improving under both scenarios. Section 5 concludes by discussing the potential research avenues opened up by our contribution.

2 SL within the canonical NK model

This section develops our reformulation of SL with respect to the canonical environment used in Arifovic et al. [2013]. Beginning with that environment, we now illustrate the steps needed to modify it to one characterized by "learning about long-term drifts with internal rationality."

2.1 The Arifovic et al. [2013] environment

We start by reproducing the SL environment used in Arifovic, Bullard, and Kostyshyna [2013]. The matrix representation of the three-equation NK model may be summarized by

$$\hat{z}_t = \alpha + B\mathbb{E}_t^* \hat{z}_{t+1} + \chi g_t, \tag{1}$$

where $\hat{z}_t = [\hat{\pi}_t, x_t]'$ is the vector of endogenous variables expressed in deviation from the deterministic steady state (the inflation gap $\hat{\pi}_t$ and the output gap x_t, respectively); $\mathbb{E}^*$ an aggregate generic expectation operator where the $*$ indicates that they need not coincide with RE; g_t a scalar representing the unique exogenous AR(1) shock; and α, B and χ coefficient matrices containing the parameters of the model. The intercept term is $\alpha = [0, 0]'$ because variables are expressed in deviations from their deterministic steady-state values, while χ is a 2-by-1 vector where the first element is 0 because g is a real shock and there is no cost-push shock.

Arifovic et al. [2013] assume a discrete population of J agents, indexed by j. These agents may differ with respect to their individual expectations, which evolve under SL and which we denote by $\mathbb{E}_{j,t}^{SL}(\hat{z}_{t+1})$ (we detail the SL algorithm in Section 2.3). Individual expectations are then aggregated across all agents using the arithmetic mean, i.e. $\mathbb{E}_t^* \hat{z}_{t+1} \equiv \mathbb{E}_t^{SL} \hat{z}_{t+1} = \frac{1}{J} \sum_j \mathbb{E}_{j,t}^{SL}(\hat{z}_{t+1})$, and inserted into the system (1).

As is usual in the learning literature, agents form their expectations using a PLM of the economy, which is akin to a forecasting model. Arifovic et al. [2013] assume that its form is identical across all agents but that agents may use different coefficient values in their respective PLMs. Additionally, the PLMs have the same functional form as the ALM, which is given by the minimum-state variable (MSV) solution to Eq. (1). In the language of the learning literature, the PLMs are well-specified. Each agent j then forms her expectations using the following PLM

$$\hat{z}_t^{(j)} = a_{j,t} + c_{j,t} g_t, \tag{2}$$

where the shock g_t is assumed to be observable at the beginning of period t and $a_{j,t}$ and $c_{j,t}$ are (2×1) vectors. Eq. (2) implies that all agents understand that the endogenous variables are (linearly) driven by the shock g around their respective steady-state values, represented by the intercepts $a_{j,t}$. They may, however, disagree on the exact values of these intercepts (and then use different values of $a_{j,t}$) *and* on the magnitude of the impact of the shock on the economy (and then use different $c_{j,t}$).

Iterating the forecasting model (2) forward gives the forecasts of each SL agent j in any period t

$$\mathbb{E}_{j,t}^{SL}(\hat{z}_{t+1}) = a_{j,t} + c_{j,t} \rho g_t, \tag{3}$$

where $\mathbb{E}_{j,t}^{SL}(\cdot)$ refers to expectations formed under SL by agent j based on her latest available coefficients $a_{j,t}$ and $c_{j,t}$. Following Arifovic et al. [2013], we have further assumed that agents have come to learn the autocorrelation of the shock ρ – for instance, by estimating its value using the past time series of g.

The REE values of the coefficients a and c may be obtained *via* the usual undetermined coefficients method. In this small-scale model, one may easily compute them as

$$\begin{aligned} \overline{a} &= [0, 0]', \\ \overline{c} &= [I - \rho B]^{-1} \chi. \end{aligned} \tag{4}$$

In the absence of SL and under RE, expectations are homogeneous and agents use the REE values $\overline{a}$ and $\overline{c}$ in their PLMs to form their forecasts. This coincides with the ALM

$$\hat{z}_t = \overline{a} + \overline{c} g_t. \tag{5}$$

By iterating (5) forward, we find:

$$\mathbb{E}_t(\hat{z}_{t+1}) = \overline{a} + \overline{c} \rho g_t, \tag{6}$$

where $\mathbb{E}_t$ denotes the RE operator conditional on the information set available in t.

Following Arifovic et al. [2013], under SL, agents are uncertain about the values of $a_{j,t}$ and $c_{j,t}$. They form beliefs about them and revise these beliefs over time based on the accuracy of the forecasts they provide. Based on this previous work, we now introduce a set of assumptions to reformulate SL expectations and obtain a closed-form solution.

2.2 A parsimonious reformulation of the Arifovic et al. [2013] learning problem

To obtain a parsimonious SL model with a general closed-form solution, we now introduce three assumptions pertaining to the information set of the SL agents into the framework of Arifovic et al. [2013]:

Assumption 1 (Internal rationality)**.** *Agents understand that SL expectations affect the macroeconomic dynamics, hence they internalize the evolution of SL beliefs over time and recognize the influence of SL dynamics over future realizations of the endogenous variables.*

Assumption 2 (Steady-state beliefs)**.** *The 'true' effects of the exogenous shocks and the lagged endogenous variables on the economy are known and common knowledge, but agents may hold heterogeneous beliefs $a_{j,t}$ about the long-run values of the variables that they have to forecast.*

Assumption 3 (Private information)**.** *Beliefs are private, such that agents cannot observe the beliefs of others: agent k does not observe any other beliefs $a_{j,t}$, $j \neq k, l$ in the population, where agent l is the tournament pair of k in a given period.*[9]

Assumption 1 implies that SL agents now also form beliefs about the aggregate SL expectations. This is akin to treating the effect of bounded rationality as an additional state variable in the ALM (5), which then looks like

$$\hat{z}_t = \overline{a} + \overline{c}g_t + \overline{d}\varphi_t, \tag{7}$$

where φ_t depicts the aggregate beliefs (in deviation from RE). The case $\varphi_t = 0$, $\forall t$ is the nested case of RE. This formulation is reminiscent of the concept of internal rationality developed by Adam and Marcet [2011], where expectations are a function of the expected changes in the PLM. In contrast to our framework, agents in Arifovic et al. [2013] do not consider the effects of their own learning on the (future) realization of the endogenous variables. They instead assume that the ALM of the economy is the one that results from RE, which boils down to assuming that the agents' perception of the effect of learning is

[9] In any period, each agent only observes the beliefs of one other agent who is randomly selected without replacement in the entire population of J agents with equal probability to be their mate in the tournament; see Section 2.3.

$\bar{d} = 0$. Therefore, our SL agents use the following PLM:

$$\hat{z}_t^{(j)} = a_{j,t} + c_{j,t} g_t + d_{j,t} \varphi_t. \tag{8}$$

We note that SL agents do not need to know how exactly φ is aggregated from the individual beliefs $a_{j,t}$ to use a PLM of the form (8).

Assumption 2 further implies that $c_{i,t} = \overline{c}$ and $d_{i,t} = \overline{d}$, $\forall j, t$ – i.e. SL agents have no uncertainty regarding the effects of the exogenous and learning shocks on the economy.[10] The PLM (8) of our SL agents therefore simplifies to:

$$\hat{z}_t^{(j)} = a_{j,t} + \overline{c} g_t + \overline{d} \varphi_t. \tag{9}$$

Eq. (9) makes explicit our reformulation of SL beliefs as 'biased rational forecasts with internal rationality' under our two key assumptions 1 and 2. Our SL agents use RE to determine the effects of the shocks and other state variables on the economy, including the effect of the aggregate deviation from RE. However, they may disagree about the long-run drifts or steady-state values of the variables that they need to forecast. They may hold different beliefs about these values, beliefs which are contained in an agent-level vector $a_{j,t}$ that evolves over time under SL (see Section 2.3). For instance, heterogeneous beliefs regarding long-run drifts in inflation may represent imperfect knowledge of the CB target – or the target's imperfect credibility.

Before pursuing the derivations, we should briefly point out the parsimony of our approach, which is visible even in this small-scale example. The learning problem of our SL agents is considerably simplified with respect to the formulation of Arifovic et al. [2013]: from four coefficients, SL agents now learn only two. This is particularly appealing if one aims to implement SL in a larger-scale DSGE model, where the number of parameters in the matrix c, the number of additional state variables (and, hence, additional matrices c) and the anticipation of the effects of learning on all these variables would render the learning problem prohibitively complex.

Iterating the forecasting model (9) of the SL agents forward, we obtain the individual SL expectation of agent j in our framework:

$$\mathbb{E}_{j,t}^{SL}(\hat{z}_{t+1}) = a_{j,t} + \overline{c} \rho g_t + \overline{d} \mathbb{E}_{j,t}^{SL}(\varphi_{t+1}), \tag{10}$$

where only $\mathbb{E}_{j,t}^{SL}(\varphi_{t+1})$ remains to be specified in order to formulate expectations under SL.

Assumption 3 implies that SL agents only have access to decentralized information. The information set of an SL agent j in any period t contains her private belief $a_{j,t}$ and

[10] In the three-equation model considered in this section, there is no lagged endogenous variables in the ALM; see Section 3 for a more general treatment that includes these.

the past realizations of all endogenous variables and the exogenous shocks, but excludes the beliefs of the $J-2$ other agents (i.e. all the other agents she has not been paired with during the tournament in period t, see Section 2.3 for detail). Furthermore, SL agents need not know how individual beliefs are aggregated into the macroeconomic dynamics; in particular, SL agents need not know that φ corresponds to the *average* beliefs across all of them (see below). Hence, Assumption 3 implies that agents only have access to their own (updated/post-tournament) beliefs from which to form expectations about the aggregate beliefs in the economy. As a result, we assume that an agent uses their idiosyncratic (post-tournament) belief $a_{j,t}$ as a proxy of the aggregate belief φ, i.e. $\mathbb{E}^{SL}_{j,t}(\varphi_{t+1}) = a_{j,t}$.[11] Individual expectations under SL therefore read as follows:

$$\mathbb{E}^{SL}_{j,t}(\hat{z}_{t+1}) = a_{j,t} + \bar{c}\rho g_t + \bar{d}a_{j,t}. \tag{11}$$

We use the arithmetic mean to aggregate the individual beliefs in Equation (11) and we refer the reader to Arifovic et al. [forthcoming, Appendix A] for the micro-foundations consistent with this aggregation procedure in the log-linearized form of the three-equation NK model.[12] Averaging Eq. (11) over the J agents yields the aggregate SL expectations in any period t:

$$\mathbb{E}^{SL}_{t}(z_{t+1}) = \frac{1}{J}\sum_{j=1}^{J}\mathbb{E}^{SL}_{j,t}(z_{t+1}) = \frac{1}{J}\sum_{j=1}^{J}a_{j,t} + \bar{c}\rho\hat{g}_t + \bar{d}\frac{1}{J}\sum_{j=1}^{J}a_{j,t}, \tag{12}$$

which may be rewritten as:

$$\mathbb{E}^{SL}_{t}(z_{t+1}) = \varphi_t + \bar{c}\rho\hat{g}_t + \bar{d}\varphi_t, \tag{13}$$

where φ_t is the aggregate deviation from RE introduced in Eq. (7) and is now explicitly expressed as the average deviation of the individual beliefs from the (zero) steady-state values, i.e. $\varphi_t \equiv \frac{1}{J}\sum_{j=1}^{J} a_{j,t}$. In other words, φ_t represents the average beliefs about the long-run value or drift of the output and inflation gaps among SL agents.

[11] This assumption is not essential. Another formulation would be to assume that SL agents also observe the aggregate belief φ_{t-1}, where we use $t-1$ because the aggregate beliefs that result from the tournament (see Section 2.3) cannot be known *ex ante* before individual expectations are formed. Furthermore, contrary to RE agents, SL agents cannot extract φ from the realizations of $\hat{z}$ using (9) because they may be uncertain about the RE values $\bar{a}$. We would need to imagine that, for instance, past aggregate beliefs φ_{t-1} are published based on the surveys conducted by policy institutions. We would further need to adjust the timing of the information set under RE so that that φ_t is not observable in t, only φ_{t-1} is. Averaging all the PLMs (10) under SL using $\mathbb{E}^{SL}_{j,t}(\varphi_{t+1}) = \varphi_{t-1}$ would result in the same aggregate form as (22) below. One may use either of the two assumptions. Note that the present version of the Dynare toolbox implements the one in the main text.

[12] Our micro-foundations hold no matter the exact form of SL used, as long as the population of heterogeneous expectations is discrete. In this companion paper, the micro-foundations also allow us to derive the second-order approximation of the welfare criterion necessary to assess optimal monetary policy as illustrated in Section 4.

We note that the heterogeneity of beliefs implies strategic uncertainty, which leaves room for higher-order beliefs. For instance, even if an agent were to know the 'true' drift values $\overline{a} = 0$, they may anticipate that others do not. This means that $\frac{1}{J}\sum_{j=1}^{J} a_{j,t} = \varphi_t \neq \overline{a}$, and agent j must adjust their own beliefs $a_{j,t}$ accordingly.

Before solving the model under SL using standard methods, it is useful to recap from the previous works aforementioned how the individual beliefs $a_{j,t}$ evolve under SL.

2.3 Unpacking the SL algorithm

There are various versions of SL in the related literature and the work of Jasmina. Both here and in the toolbox we will use a parsimonious formulation involving only three steps.

In the *first step*, at the beginning of each period t, some agents receive news; this is the exploration phase of the strategy space. We denote the news received by agent j by $m_{j,t}$ (for *m*utation in the GA language). Mutations concerning each strategy (i.e. each value in the vector $a_{j,t}$ used in the PLM 9) occur with probabilities collected in a vector μ.[13] Given a large number of agents, in each period each agent receives news about each variable, with a given probability contained in μ. These pieces of news are white noise shocks that perturb the existing beliefs contained in $a_{j,t-1}$. Formally, after the news has been received, each agent's 'mutated' beliefs are given by:

$$m_{j,t} = a_{j,t-1} + \mathbb{1}_{\varpi_{j,t}\leq\mu}\,\iota_{j,t}\xi. \tag{14}$$

Vector $\varpi_{j,t}$ contains a series of independent random draws from a standard normal distribution, while the vector $\mathbb{1}_{\varpi_{j,t}\leq\mu}$ is a vector of dummy variables where each entry is equal to one if the corresponding draw in $\varpi_{j,t}$ is lower than the corresponding probability value in vector μ, and zero otherwise. In other words, the dummy variables indicate whether the agent has received a news about a given variable or not. It is therefore possible (likely) that only some agents receive news about some –but not all– variables in any period. Vector $\iota_{j,t}\xi$ is the vector of white-noise shocks representing the news, drawn from standard normal distributions (where the draws are gathered in the vector $\iota_{j,t}$); it is scaled by ξ the covariance matrix of news with a size corresponding to the number of variables subject to SL. In what follows, we assume that the news is uncorrelated across variables, such that ξ is a diagonal matrix where each diagonal element corresponds to the standard deviation of the draw associated to the news regarding a given variable.

These pieces of news constitute new candidates for the long-run values to be used in the forecasting rules (11). Before using them to form SL forecasts, however, agents must

[13] μ is a column vector of probabilities of size corresponding to the number of variables to forecast under SL (i.e. the number of non-zero values in the vector $a_{j,t}$; see Footnote 19). The frequency of news arrival needs not be the same for all these variables but we assume that these probabilities, as the standard deviations of the mutations, are constant over time and across agents. Each row of μ corresponds to the probability of receiving a news for the corresponding variable in the vector $a_{j,t}$.

filter out the least accurate or relevant ones.

Therefore, the *second step* of SL involves computing this accuracy. The accuracy of each updated belief $m_{j,t}$ is evaluated based on a discounted sum of past forecast errors that agent j would have made over the recent history, had she used the particular $m_{j,t}$ values given in Eq. (14) in her forecasting model. We denote by $F_{j,t}$ the vector collecting these accuracy measurements (for Fitness), which are computed as the discounted sums of the squared differences between the past realizations and the forecasts obtained by using the values $m_{j,t}$ in the PLM (11)

$$F_{j,t} = -\sum_{\tau=0}^{t-1} \delta^{\tau} \left(\hat{z}_{t-1-\tau} - \mathbb{E}^{SL}_{j,t-2-\tau} \left(\hat{z}_{t-1-\tau} | m_{j,t}\right)\right)^2, \tag{15}$$

where δ is a diagonal matrix (of size equal to the number of variables forecast using SL) with diagonal terms between 0 and 1. These terms correspond to the discount weights associated with the history of each of these variables (the history is assumed to start at $t = 0$). Note that $\mathbb{E}_{j,t-2-\tau}\left(\hat{z}_{t-1-\tau}|m_{j,t}\right)$ is the forecast of $\hat{z}_{t-1-\tau}$ that would have resulted from the forecasting model (11) with $m_{j,t}$, i.e. $\mathbb{E}_{j,t-2-\tau}\left(\hat{z}_{t-1-\tau}|m_{j,t}\right) = m_{j,t} + \bar{c}\rho g_{t-2-\tau} + \bar{d} m_{j,t}$. Eq. (15) results in a vector of forecasting accuracy for each agent j, where each entry corresponds to the accuracy of her corresponding updated beliefs in the vector $m_{j,t}$.

A short digression helps illustrate the strategic complementarity under SL. To see it, we simply replace the past values of $\hat{z}$ in (15) with the MSV expression (7) and group terms to obtain

$$F_{j,t} = -\sum_{\tau=0}^{t-1} \delta^{\tau} [\bar{c}(g_{t-1-\tau} - \rho g_{t-2-\tau}) - m_{j,t}] + [\bar{d}(\varphi_{t-1-\tau} - m_{j,t})])^2. \tag{16}$$

The most accurate mutations (or news) $m_{j,t}$, which are more likely to survive the tournament and be used to form SL expectations, are those that minimize the sum in (16). This requires the beliefs $m_{j,t}$ to be relatively close to the recent (unexpected) fundamental developments that stem from the innovation in g (in the first term within squared brackets), but also as close as possible to the aggregate beliefs $\varphi_{t-1-\tau}$ (in the second term in squared brackets). Because φ consists of the average of the individual beliefs, SL further resembles a beauty-contest game, in which better performing strategies are those that are as close as possible to the population average. In other words, agents have an interest in guessing average beliefs as accurately as possible.

In the *third and final step*, the beliefs that result in lower forecast errors over the recent past (i.e. the ones with relatively high fitness values) are adopted by the agents, whereas those with lower fitness values are discarded. This is the selection phase of SL, which we implement via a tournament:[14] in every period, each agent is paired with another

[14]This implementation has the advantage of being decentralized and, therefore, consistent with the

one, where pairs are formed by random uniform draws without replacement. This process results in $J/2$ couples (where J is conveniently chosen to be even). For each pair of agents k and l, $k \neq l$, the fitness values assigned to each row of the belief vectors $m_{k,t}$ and $m_{l,t}$ are compared. Agents adopt the one with the highest fitness (i.e., the highest forecasting accuracy) and implement it in their respective PLMs – therefore, both agents end up with the same set of beliefs after the tournament. This process delivers the actual individual SL forecasts, which are then aggregated using the average in Eq. (12). Formally, the tournament phase may be written as

$$(a_{k,t}, a_{l,t}) = \mathbb{1}_{F_{k,t}>F_{l,t}}(m_{k,t}, m_{k,t}) + (1 - \mathbb{1}_{F_{k,t}>F_{l,t}})(m_{l,t}, m_{l,t}), \tag{17}$$

where the vector $\mathbb{1}_{F_{k,t}>F_{l,t}}$ is a vector of dummy variables equal to one if the corresponding belief of agent k is more accurate (has a higher fitness) than the one of agent l and 0 otherwise. We note that the fitness comparison is performed for each variable (i.e., each row of vectors m); as such, the best forecaster out of the pair need not be the same for each variable.

Taken together, it is convenient to summarize the non-linear belief updating process under SL (14)-(17) as a function $s(\cdot)$ that intuitively describes the change in beliefs of agent j in period t, i.e. which mutations occur and survive the tournament:

$$a_{j,t} - a_{j,t-1} = s\left(\iota_{j,t}\xi, \hat{z}_{t-1}, \ldots, \hat{z}_1\right). \tag{18}$$

One may find it useful to rewrite Eq. (18) in aggregate terms:

$$\begin{aligned} \varphi_t &= \varphi_{t-1} + s\left(\iota_t, \hat{z}_{t-1}, ..., \hat{z}_0\right) U \\ &\equiv \varphi_{t-1} + S_t, \end{aligned} \tag{19}$$

where U is a vector of aggregation with values of $1/J$ and S_t is the compact update step that depends on both the aggregation of individual forecasts and the selected news following the tournament.

We now have the necessary ingredients for solving the SL model within the canonical three-equation model using perturbation methods.

2.4 Solving the reformulated SL model with standard methods

To use standard solution methods to obtain a closed-form solution, we need two last steps: i) defining the concept of RE within the SL model, and ii) rewriting the model as a function of RE.

information assumption 3 that states that SL agents do not observe the beliefs of all other agents.

First, one can think of RE from the point of view of an agent overseeing the SL economy who is aware of the presence of (boundedly rational) SL agents.[15] An RE agent knows the true model of the economy (i.e., she knows the ALM (7) and, in particular, that $\bar{a} = 0$) and is thus able to infer the aggregate bias φ.

After iterating the ALM (7) forward and taking mathematical expectations, RE in our framework are given by

$$\mathbb{E}_t(\hat{z}_{t+1}) = \bar{c}\rho g_t + \bar{d}\mathbb{E}_t(\varphi_{t+1}), \tag{20}$$

where we have set $\bar{a} = 0$ (which holds in any log-linearized model). The term $\mathbb{E}_t(\varphi_{t+1})$ makes it clear that RE agents need to anticipate the outcomes of SL – the outcome of the selection of the news given by (18) – in the next period. Per Eq. (19), we can write $\mathbb{E}_t(\varphi_{t+1}) = \varphi_t + E_t(S_{t+1})$; therefore, the RE agent needs to anticipate the change in the population of beliefs in the next period (S_{t+1}). We make the following assumption:[16]

Assumption 4 (Imperfect information under RE)**.** *An RE agent cannot observe the individual beliefs (possibly updated by news) of the SL agents because such beliefs are private information.*

Whenever individual variables are unobservable, even an RE agent cannot predict how the average population of SL beliefs φ will be modified in the next period.[17] Therefore, the best she can do is to use the latest observable information about the bias of aggregate expectations with respect to RE. In other words, the aggregate bias in beliefs φ is treated as a random-walk process. We therefore can write $\mathbb{E}_t(\varphi_{t+1}) = \varphi_t$, which results in

$$\mathbb{E}_t(\hat{z}_{t+1}) = \bar{c}\rho g_t + \bar{d}\varphi_t. \tag{21}$$

From there, in order to solve the SL model with standard techniques, we only need to make use of the fact that the aggregate SL expectations (13) may be expressed as a

[15] Our RE agent is therefore different from the so-called ‘fundamentalist agents’ usually considered in the heterogeneous-expectation literature who form expectations in line with the steady-state path that would prevail in the absence of learning. As shown in this section, our RE agent instead correctly anticipates (on average and given her information set) the effects of boundedly rational expectations on the future economy.

[16] We remark that our formulation of the RE counterpart is reminiscent of dispersed information models à la Lorenzoni [2010].

[17] In the absence of this information friction, an RE observer would be able to compute the expected fitness of each updated beliefs $m_{j,t}$ based on an expected realization of $\hat{z}_t$ and, therefore, an expected fitness value of each updated belief in $t+1$, and then anticipate (on average) the tournament output from all the potential $J(J-1)/2$ combinations of pairs to compute (in expectation) the change in the aggregate beliefs S_{t+1}. In other words, we assume that an RE agent does not have access to enough information to do so because she cannot observe individual variables and, hence, treat mutation and tournament shocks as *ex-ante* white noise.

function of RE (21):

$$\mathbb{E}_t^{SL}(\hat{z}_{t+1}) = \varphi_t + \mathbb{E}_t(\hat{z}_{t+1}), \tag{22}$$

where we see that φ_t is treated as an additional state variable that exists under SL and represents the average deviations from RE introduced by SL.

The expectations E_t^* in the reduced-form model (1) may be replaced by (22), giving us

$$\hat{z}_t = \alpha + B(\mathbb{E}_t\hat{z}_{t+1} + \varphi_t) + \chi g_t. \tag{23}$$

Note that Eq. (23) resembles any DSGE model in reduced form and may be solved by the standard methods implemented in Dynare. To do so, we can substitute out the RE term in (23) using (21). Rearranging the terms yields the following (omitting the bars for clarity):

$$\hat{z}_t = \alpha + (Bc\rho + \chi)g_t + (Bd + B)\varphi_t. \tag{24}$$

Identifying the terms in the MSV solution (7) and Eq. (24) solves for the SL equilibrium: $\alpha = a$, $(Bc\rho + \chi) = c$ and $Bd + B = d$. The solutions $a = \bar{a} = 0$ and $c = \bar{c}$ are the same as under RE given in Eq. (4). SL introduces an additional term pertaining to the effect of the average deviation from RE $\overline{d} = B(I - B)^{-1}$. This is the solution that our Dynare toolbox computes. Note that our formulation implies a time-invariant policy function; hence, the law of iterated projections is satisfied in our framework, something which is a common concern when using learning-driven expectation formation processes in reduced-form log-linearized models.

Having outlined our reformulation of the small-scale model used in Arifovic et al. [2013], we now provide a general formulation for embedding SL expectations into any DSGE model and introduce the main functions of the Dynare toolbox.

3 A general formulation of DSGE models with SL expectations

To set the stage for the introduction of the toolbox, we will now provide textbook derivations of an otherwise standard DSGE model with SL expectations in a self-contained presentation by using the notation and model structure of Dynare.

3.1 A general solution for SL models using perturbation methods

A dynamic model with both forward- and backward-looking variables can be cast in a state-space form as follows:

$$\mathbb{E}_t \left(f_\Theta \left(z_{t+1}, z_t, z_{t-1}, \varepsilon_t, \mathbf{a}_t \right) \right) = 0, \tag{25}$$

where f_Θ is the vector of $N \times 1$ non-linear equations with structural parameters stacked into vector Θ, the term $\mathbb{E}_t(\cdot)$ denotes RE based on the information set available at t, z_t is the vector of $N \times 1$ endogenous variables in Dynare notation,[18] ε_t is the vector of $N_\varepsilon \times 1$ stochastic innovations $\varepsilon_t \sim \mathcal{N}(0, \Sigma_\varepsilon)$ (i.e. the exogenous shocks), and $\mathbf{a}_t$ is a matrix of size $J \times N$ gathering the idiosyncratic beliefs $a_{j,t}$ for the N variables of the J agents at time t.[19]

In our setup, the SL idiosyncratic beliefs $\mathbf{a}_t$ only affect the forward-looking variables, which allows us to write the expected vector as a function $g(z_{t+1}, \mathbf{a}_t)$. The simplified system reads as

$$\mathbb{E}_t \left(f_\Theta \left(g \left(z_{t+1}, \mathbf{a}_t \right), z_t, z_{t-1}, \varepsilon_t \right) \right) = 0. \tag{26}$$

A first-order Taylor expansion of function $f_\Theta(\cdot)$ around a fixed point is given by

$$\mathbb{E}_t \left(F \left(D_1 \hat{z}_{t+1} + D_2 \varphi_t \right) + G \hat{z}_t + H \hat{z}_{t-1} + M \varepsilon_t \right) = 0, \tag{27}$$

Matrices F, G and H are the $N \times N$ Jacobian matrices of $f_\Theta(\cdot)$ in z_{t+1}, z_t and z_{t-1} respectively. M is the $N \times N_\varepsilon$ Jacobian matrix of system f_Θ in the vector of stochastic (fundamental) innovations ε_t. φ_t results from the log-linearization of $g(\cdot)$ which gives SL expectations as a weighted sum of the RE component $\hat{z}_{t+1}$ and the vector of the average deviations of the idiosyncratic beliefs from their zero steady-state values φ_t (the weights D_1 and D_2 result from the log-linearization of the function g). Moreover, we can define $\varphi_t \equiv \mathbf{a}_t U$, where U is an aggregation vector defined in (19) and φ_t is an $N \times 1$ vector in which only $N_{SL} \times 1$ items are non-zero. Furthermore, $\hat{z}_t = z_t^{(1)} - \bar{z}$ is the first-order Taylor expansion term around the fixed point $\bar{z}$. As our expansion includes only first-order terms, $z_t^{(1)}$ is the linear-implied path consistent with the expansion order considered. Note that Eq. (27) is a generalization of the reduced-form three-equation model (23) that allows

[18] In a standard Dynare file, all endogenous variables are stacked into a VAR(1) process. Using this standard notation in Section 2, the AR(1) shock g_t is endogenous and contained in $\hat{z}$ and its i.i.d. component is exogenous and contained in ε.

[19] Under this formulation, agents forecast a subset $N_f \leq N$ of these variables, of which one may specify $N_{SL} \leq N_f$ that are subject to SL, while the remaining $N_f - N_{SL}$ are subject to RE. Our toolbox allows the user to freely declare which variables are subject to SL or RE forecasts. The $N - N_{SL}$ columns of $\mathbf{a}_t$ corresponding to the non-SL variables are then simply filled with zeros while the remaining N_{SL} columns contain the population of beliefs $\{a_{j,t}\}$ that evolve under SL. For instance, in this paper, we assume that the AR(1) shocks are not subject to SL.

for lagged endogenous variables, as well as any number of exogenous shocks and state variables contained in vectors z and ϵ respectively.

Omitting the zero-vector of intercepts, the general MSV solution of a log-linearized system (27) has the following form:

$$\hat{z}_t = P\hat{z}_{t-1} + Q\varepsilon_t + R\varphi_t, \tag{28}$$

where P, Q and R are three unknown matrices that can be solved for using standard perturbation methods embedded in our Dynare toolbox.

To do so, we first need to express the SL beliefs as a function of RE in order to identify matrices D_1 and D_2 in Eq. (27). In the vocabulary of the learning literature, Eq. (28) corresponds to the ALM, from which we can derive the PLM of agent j, i.e. her forecasting model under SL:

$$\hat{z}_t^{(j)} = a_{j,t} + P\hat{z}_{t-1} + Q\varepsilon_t + R\varphi_t. \tag{29}$$

Applying this model to both the nowcast and the forecast (by a forward iteration) of $\hat{z}$ gives the individual SL beliefs of agent j:

$$\begin{aligned}\mathbb{E}_{j,t}^{SL}(\hat{z}_{t+1}) =& a_{j,t} + P\mathbb{E}_{j,t}^{SL}\hat{z}_t + R\mathbb{E}_{j,t}^{SL}\varphi_{t+1},\\ =&(I+P)a_{j,t} + P^2\hat{z}_{t-1} + PQ\varepsilon_t + (PR+R)a_{j,t},\end{aligned} \tag{30}$$

where we have set

$$\mathbb{E}_{j,t}^{SL}\left(\varepsilon_{t+1}\right) = \mathbb{E}_t\left(\varepsilon_{t+1}\right) = 0, \tag{31}$$

$$\mathbb{E}_{j,t}^{SL}\left(\varphi_{t+1}\right) = \mathbb{E}_{j,t}^{SL}\left(\varphi_t\right) = a_{j,t}, \tag{32}$$

which is consistent with the zero mean distribution of aggregate shocks and Assumption 3 above. Eq. (30) is the counterpart of Eq. (11) in the general case.

Averaging Eq. (30) across the J agents gives the aggregate SL expectations:

$$\mathbb{E}_t^{SL}(\hat{z}_{t+1}) = (I+P)\varphi_t + P^2\hat{z}_{t-1} + PQ\varepsilon_t + (PR+R)\varphi_t. \tag{33}$$

We now compute RE. In line with Assumption 4, we know that the expected aggregate beliefs under RE read as

$$\mathbb{E}_t\left(\varphi_{t+1}\right) = \varphi_t. \tag{34}$$

Iterating (28) forward, taking the mathematical expectations (while assuming that $\hat{z}_t$

is not part of the information set in t) gives the RE:

$$\mathbb{E}_t\left(\hat{z}_{t+1}\right) = P^2\hat{z}_{t-1} + PQ\varepsilon_t + (PR + R)\varphi_t. \tag{35}$$

Accordingly, the SL expectations (33) may also be written as a function of RE:

$$\mathbb{E}_t^{SL}(\hat{z}_{t+1}) = \mathbb{E}_t(\hat{z}_{t+1}) + (I + P)\varphi_t, \tag{36}$$

where we identify $D_1 \equiv I$ and $D_2 \equiv I + P$ in Eq. (27).[20] We may now solve for matrices P, Q and R using standard methods. To do so, substitute (35) into (27), rearrange the terms and identify the coefficients (see Appendix A for details). A policy function solving (27) must satisfy these three conditions:

$$\begin{aligned} P &= -\left(FP + G\right)^{-1} H, \\ Q &= -\left(FP + G\right)^{-1} M, \\ R &= -\left(FP + G\right)^{-1} F\left(R + D_2\right), \text{ with } D_2 = P + I. \end{aligned} \tag{37}$$

Matrices P and Q are the same as in the RE model (computed already by Dynare) but the SL dynamics involve an additional matrix R that our toolbox solves for.

Because P admits two solutions, there are also two corresponding matrices Q and two matrices R. Following the usual practice in the macroeconomic literature, we only consider the stable root of the problem to simulate our model. Therefore, the RE model needs to be determined before the toolbox can simulate its SL counterpart.[21] The toolbox is therefore particularly handy for larger-scale models than the three-equation example, as the toolbox relies on the perturbation methods implemented in Dynare to quickly and accurately find the solution matrices P, Q and R.

Before turning to the application of the toolbox, it is useful to discuss the stability of the REE under our formulation of SL.

3.2 Determinacy under RE and stability under SL

Matrices P, Q, and R in Eq. (37) do not depend on the SL developments given by Eq. (19). The determinacy conditions under RE – namely, the parameter restrictions $\max|\lambda_P| \in \mathbb{C} < 1$ that ensure that the eigenvalues of the P matrix lie within the unit circle – need to be satisfied in order to use the model under SL, as explained above. To further investigate stability under SL, one may rewrite the policy function in a recursive

[20] In the absence of an endogenous lagged forecast variable, $P = 0$ and $D_2 = I$. This is in line with Eq. (22) in Section 2.

[21] In the toolbox, the function `get_PQ` first solves for the RE solution and selects the solution P that ensures determinacy.

way using (19):

$$\hat{z}_t = P\hat{z}_{t-1} + Q\varepsilon_t + R\varphi_{t-1} + RS_t.$$

The complete dynamical system reads as:

$$\begin{bmatrix} \varphi_t \\ \hat{z}_t \end{bmatrix} = \begin{bmatrix} I_{N_y} & 0_{N_y} \\ R & P \end{bmatrix} \begin{bmatrix} \varphi_{t-1} \\ \hat{z}_{t-1} \end{bmatrix} + \begin{bmatrix} 1_{N_y} & 0_{N_y} \\ R & Q \end{bmatrix} \begin{bmatrix} S_t \\ \varepsilon_t \end{bmatrix}. \tag{38}$$

We therefore obtain the following partial result:

Proposition 1. *Determinacy under RE is a necessary but not sufficient condition for stability under SL.*

Proof. The system (38) is stable under SL iif all eigenvalues of matrix $\Omega \equiv \begin{bmatrix} I_N & 0_N \\ R & P \end{bmatrix}$ lie within the unit circle. Matrix Ω is a triangular block matrix. Hence, we have $det(\Omega) = det(I_N)det(P)$, or $det(\Omega - \lambda I_N) = det(1 - \lambda I_N)det(P - \lambda I_N)$, where the roots $\lambda \in \mathbb{C}$ are the eigenvalues of Ω. It follows that the eigenvalues of Ω are the eigenvalues of P – which lie within the unit circle iif the model (28) is determinate under RE – and 1 (with multiplicity N). Determinacy under RE is not a sufficient condition for stability under SL because the presence of eigenvalues on the unit circle (independently of the parameter values of the model) renders the case non-generic. □

In other words, stability of the REE under SL may or may not result if the model is determinate and depends on the non-linearities introduced by the SL expectations *via* the function $s(\cdot)$, for which no analytical solution is available.[22]

One of the contributions of this paper is the incorporation of the SL solution and simulation methods into MATLAB [Inc., 2022] as a toolbox of the Dynare package [Juillard et al., 1996]. We will now briefly detail how to use it.

3.3 Overview of the Dynare toolbox for SL models

Our SL toolbox is seamlessly implemented using standard Dynare syntax. Users may write their SL model in a `.mod` file as they otherwise would with an RE model. They first need to define which variables will be forecast using SL (the other forward variables will then be forecast using RE). They then need to specify the values of the structural parameters of the model and of the SL algorithm and the sequences of the exogenous shocks. They may also choose whether the SL news shocks (the mutations) are drawn from a predefined seed or not. The only remaining requirement for solving and simulating a model under SL is to first call the function `SL_get_pq`, followed by `SL_chain`. This is

[22]Intuitively, these eigenvalues on the unit circle stem from the random-walk beliefs about the evolution of SL.

necessary because the function `SL_get_pq` solves for matrices P, Q and R (see Eq. (37)), after which the function `SL_chain` then simulates the model under RE and SL, given these matrices and the chain of exogenous shocks provided. The simulations under RE are produced for the sake of comparison with the model under SL. Our toolbox allows one to simulate both IRFs and stochastic time series of the model, as is common in standard RE macroeconomic literature.

The toolbox files provided alongside this paper[23] contain extensive comments. We have also included the code to implement the model from Section 2 and derive the results from Section 2.4. For illustration purposes only, Appendix B displays an example of a run and Appendix C reports IRFs given a shock to the IS curve in the three-equation model used in the application in Section 4.1 below. The IRFs include the responses of the cross-sectional dispersion in individual expectations over time.

4 Application of the Dynare toolbox

We illustrate the added value of the micro-founded SL framework, our solution method and the associated `Dynare` toolbox with an application to the design of optimal monetary policy rules under SL and RE. We first discuss the baseline model used.

4.1 A simple framework under SL

In [Arifovic et al., forthcoming, Appendix A] and Grimaud et al. [2024], we developed the micro-foundations of a baseline NK model with history-dependent monetary policy. The backbones of the model read as follows:

$$\begin{aligned}
\hat{y}_t &= \mathbb{E}_t^*(\hat{y}_{t+1}) - \frac{1}{\sigma}\left(\hat{\imath}_t - \mathbb{E}_t^*\left(\hat{\pi}_{t+1}\right)\right) + g_t,\\
\hat{\pi}_t &= \kappa\hat{y}_t + \beta\mathbb{E}_t^*\hat{\pi}_{t+1} + u_t,\\
\hat{\imath}_t &= \phi_\pi\hat{\tau}_t + \phi_y\hat{y}_t\\
\hat{\tau}_t &= \hat{\pi}_t + (1-\omega)\hat{\tau}_{t-1},\\
g_t &= \rho_g g_{t-1} + \varepsilon_t^g,\\
u_t &= \rho_u u_{t-1} + \varepsilon_t^u,
\end{aligned}$$

where the notations correspond to those used in Section 2 and the equations describe, by order, an IS curve, a NK Phillips curve, a Taylor rule, a weighted average of the past inflation gaps, a real shock process and a cost-push shock process denoted by u. The only peculiarity of this model is the Taylor rule that sets the nominal interest rate $\hat{\imath}$ (expressed

[23] We will provide the toolbox using `mendeley` (a free and open-access data repository) which will provide a DOI and refer to it here.

in deviation from its steady-state value) as a function of the past average inflation gap $\hat{\tau}$, which corresponds to a history-dependent inflation target (see, e.g., Budianto et al. 2023). This formulation presents the advantage of nesting the two polar cases which correspond to $\omega = 1 \Leftrightarrow \hat{\tau}_t = \hat{\pi}_t$ or inflation targeting (IT, hereafter) and price-level targeting (PLT, hereafter), when $\omega = 0$ and $\hat{\tau}_t = \hat{\pi}_t + \hat{\tau}_{t-1} = \hat{\pi}_t + \hat{P}_{t-1} = \hat{P}_t$, i.e. the price gap with respect to the path growing at the targeted inflation rate. For any intermediary values of ω, the CB implements average-inflation targeting (AIT, hereafter) with a geometric discounting of past inflation gaps. Recall that the notation $\mathbb{E}_t^*$ accommodates either RE $\mathbb{E}_t$ or aggregate SL expectations $\mathbb{E}_t^{SL}$.

Additionally, the microfoundations allow us to derive a welfare criterion for the analysis of optimal policy rules for a three-equation NK model. In the aforementioned contributions, we show that the average utility of an SL agent j, which we use as a measurement of aggregate welfare, denoted by $\mathcal{W}_t$, may be decomposed as:

$$\mathcal{W}_t = E\left(U_{j,t}\right) \simeq u_0 - u_{\gamma\gamma} Evar_t\left(\hat{\gamma}_{j,t}\right) - u_{yy} var_t\left(\hat{y}_t\right) - u_{\rho\rho} Evar_t\left(\hat{\rho}_{j,t}\right) - u_{\pi\pi} Evar_t\left(\hat{\pi}_{j,t}\right), \quad (39)$$

where the operator E is the probabilistic mean, u_0 is the steady-state level of welfare, $\hat{\gamma}_{j,t} = \hat{c}_{j,t} - \hat{C}_t$ is the percentage deviation of the consumption of agent j from aggregate consumption C_t, $\hat{\rho}_{j,t} = \hat{p}_{j,t} - \hat{P}_t$ is the relative price of firm j and $u_{\gamma\gamma}, u_{yy}, u_{\rho\rho}, u_{\pi\pi} > 0$ are the elasticities of the welfare function with respect to the variances of, respectively, the idiosyncratic consumption levels, the output gap, the idiosyncratic price levels and the idiosyncratic inflation rates. The toolbox function `get_idio` generates the idiosyncratic variables that are necessary to compute the welfare, the function `get_welff` then computes the corresponding welfare through Eq. (39).

4.2 Calibration

We use standard parameter values for the fundamental parameters of the model because the purpose of the exercise hereafter is to illustrate the use of the toolbox rather than to fit an empirical dataset. The discount factor is set to $\beta = 0.99$. $\sigma = 1$ implies a log utility function in the separable utility case, which corresponds to $\sigma' = 2.5$ in the non-separable case (see Grimaud et al. [2024]). The elasticity of substitution is set to $\varepsilon = 4$ (or, equivalently, the markup is set at 33%), and the slope of the Phillips curve to $\kappa = 0.05$. With an inverse of the Frish elasticity equal to $\varphi = 1$, these parameter values correspond to Rotemberg adjustment costs of $\xi' = 59.7015$ in the non-separable case. The benchmark monetary policy parameters represent a standard IT rule à la Taylor [1993] without interest-rate smoothing $\{\phi_\pi, \phi_y, \omega\} = \{1.5, 0.125, 1\}$. This parameter set fulfills the Taylor principle and ensures the determinacy of the baseline model under

RE.[24] The yearly inflation target is set to 2%. As for the fundamental shocks, we use $\{\rho_g, \rho_u\} = \{0.8, 0.5\}$, $\sigma_\varepsilon^g = 0.1$ and $\sigma_\varepsilon^g = 0.05$.

The SL process also involves some parameters (see, again, Section 2.3 for the notations). We follow, *inter alia*, Arifovic et al. [2018, forthcoming] and set the mutation probabilities (which correspond to the frequencies of news reception regarding the corresponding variable) to $\mu = (\mu_\pi, \mu_y)' = (0.3, 0.3)'$ and the memory in the discounted fitness function (15) to $\delta = (\delta_\pi, \delta_y)' = (0.8, 0.8)'$. The variance of the news shocks in Eq. (14) is set to $\xi = \begin{pmatrix} \sigma_{\xi,\pi}^2 & 0 \\ 0 & \sigma_{\xi,y}^2 \end{pmatrix} = \begin{pmatrix} 0.3^2 & 0 \\ 0 & 0.3^2 \end{pmatrix}$ i.e. news shocks have the same normal distribution with a standard deviation of 0.3 percent for both the inflation and output gaps. Finally, we use $J = 300$ agents.

4.3 Stabilization trade-off under IT and SL

The toolbox makes it easy to contrast the stabilization trade-off for monetary policy under RE and SL. In this section, we consider simple IT policies and set $\omega = 1$, which reduces the monetary policy rule to a standard IT Taylor rule with two parameters ϕ_π, ϕ_y, namely the reaction coefficients to the inflation and the output gaps. We decompose the welfare analysis under RE and SL as follows: we first look into how monetary policy may stabilize the aggregate components of welfare, then the idiosyncratic components that arise under SL, and conclude with the welfare implications of SL with respect to the RE benchmark.

4.3.1 Aggregate variables

In the textbook three-equation model that we consider here, the stabilization policy under RE is a trade-off between stabilizing the variances of inflation and output. In what follows, we analyze the variance of inflation and output as a function of parameters ϕ_π, ϕ_y. We set $\phi_\pi \in (1, 10]$ and $\phi_y \in [0, 10]$ to ensure determinacy of the RE model. Fig. 1 illustrates the variances of the aggregate variables, namely inflation (top panel) and the output gap (bottom panel) under RE (left-hand side panel) and SL (right-hand side panel) as a function of these two IT parameters.

Both under RE and under SL, the results are in line with the textbook presentation, where the volatility of the inflation gap is a decreasing function of ϕ_π (see Figs. 1a and 1b) and the volatility of the output gap decreases with ϕ_y (see Figs. 1c and 1d). However, both the volatility of inflation and of output are higher under SL than under RE, with the contrast being most striking regarding inflation volatility.

Fig. 2 provides another look at this trade-off. It displays the so-called IT efficiency frontier under RE (solid line) and under SL (dashed line), which is the constellation

[24] In Subsection 4.3 where we explore the monetary policy parameter space, we only consider values that ensure a determinate model under RE.

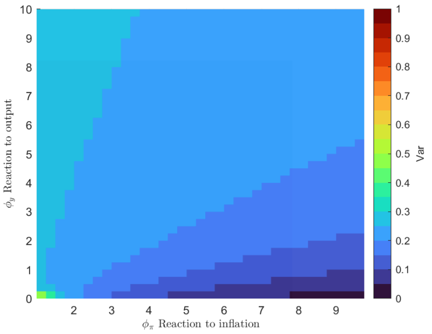

(a) Inflation variance under RE $var(\hat{\pi}_t)$

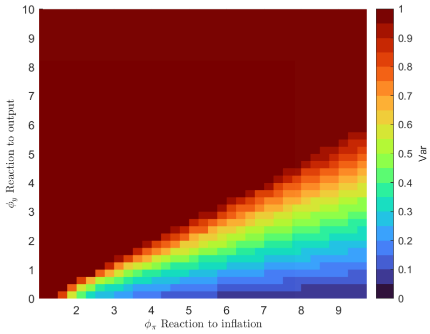

(b) Inflation variance under SL $var(\hat{\pi}_t)$

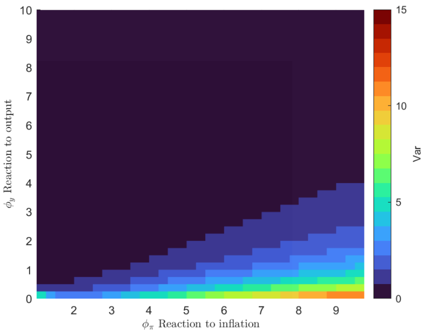

(c) Output variance under RE $var(\hat{y}_t)$

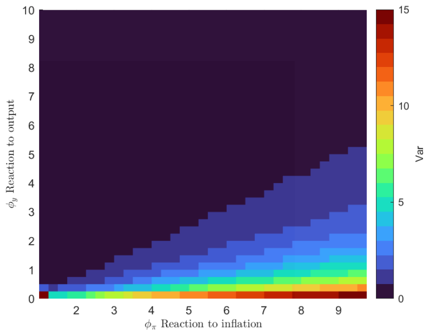

(d) Output variance under SL $var(\hat{y}_t)$

Notes: We use a discrete grid of the value for the two reaction coefficients in the standard IT policy rule. For each combination, the moments are computed based on the average value over 50 chains of 1,000 aggregate shocks that are taken to be the same under RE and SL. The random draws of the SL algorithm are independently drawn for every simulation. Each time series includes 1,000 periods excluding a 300-period burn-in phase. The change in color is truncated in order to improve readability.

Figure 1: Trade-off with respect to the stabilization of aggregate variables

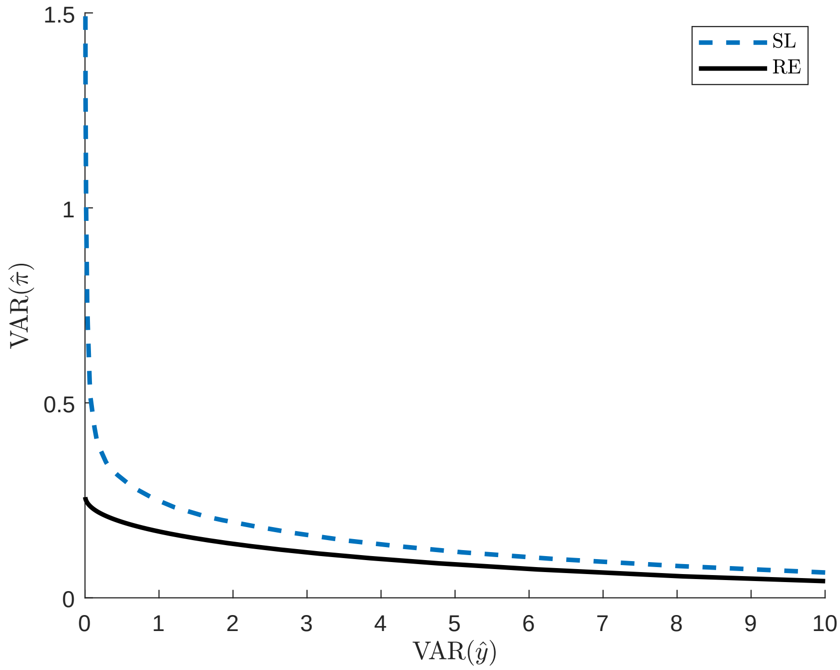

Figure 2: Efficiency Frontiers

of points that represent the minimum inflation variance (on the x-axis) and output gap variance (on the y-axis) obtained for each weight $\alpha \in [0, 1]$ in the following loss function:[25]

$$\mathcal{L}(\alpha) = \alpha Var(\hat{\pi}_t) + (1 - \alpha) Var(\hat{y}_t). \tag{40}$$

The closer the frontier to the origin, the more efficient monetary policy because it entails a lower aggregate volatility than policies on frontiers further up (which corresponds to a higher inflation volatility) or further to the right (which corresponds to a higher output gap volatility).

The main takeaway from Fig. 2 is that monetary policy is more efficient under RE than under SL, in particular regarding inflation stabilization: the SL frontier is clearly further up the RE frontier and the two do not intersect, which means that there is no combination of monetary policy parameters that could deliver the relatively lower aggregate volatility obtained under RE when agents have SL expectations.

In detail, the RE and the SL frontiers further diverge from each other on the right-hand side of the plot, when the weight on inflation stabilization tends towards 1. While the variance of inflation converges towards 0 under RE, it converges towards a positive asymptote under SL. This divergence of outcomes under the two expectation schemes is due to the higher volatility in inflation expectations, and hence in inflation, under SL than under RE as a result of the idiosyncratic news. This excess volatility particularly

[25]Hence, the minimization problem reads as $\min_{\{\phi_\pi, \phi_y\}} \mathcal{L}(\alpha), \forall \alpha \in [0, 1]$.

affects inflation, more than the output gap, because inflation is almost self-fulfilling in the NK model (β is near unity in the NK Phillips curve, while its slope κ has a much smaller value).

An additional major difference between RE and SL arises when the CB prioritizes stabilizing output (on the left-hand side of Fig. 2). In this region, the SL frontier is steeper than the RE frontier and the two diverge strongly from each other. This means that under SL, drastically reducing output volatility comes at a high cost in terms of inflation volatility. In other words, the so-called sacrifice ratio in the SL model is higher than in the RE model. This is a well-known result for NK model under boundedly rational expectations such as under adaptive learning [Orphanides and Williams, 2005].

4.3.2 Trade-off with respect to the stabilization of the idiosyncratic components of welfare

In Fig. 3, we illustrate the variances of the idiosyncratic components of the welfare function (39), namely the dispersion of the relative prices $\hat{\rho}_{j,t}$ (top panel) and consumption levels $\hat{\gamma}_{j,t}$ (bottom panel) under RE (left-hand side panel) and under SL (right-hand side panel) as a function of the two IT parameters ϕ_π (x-axis) and ϕ_y (y-axis).

Of course, by construction, the RE values of these idiosyncratic welfare components are always zero but we report them for the sake of comparison with SL. By contrast, under SL, heterogeneity implies additional volatility besides the variance of the aggregate variables. In particular, the dispersion of the individual relative prices ($\hat{\rho}_{j,t}$) co-moves with the variance of aggregate inflation; to see that superpose Figs. 1b and 3b. This is because the more volatile inflation, the more disagreement between agents and the more dispersed their prices. Intuitively, more volatile inflation implies that the mutations of SL may generate non-steady state expectations that temporarily survive the tournament because they happen to generate the lowest forecast errors over a recent volatile history and, then, turn self-fulfilling. More volatile inflation then implies larger and more frequent shifts in the inflation trend. Along these shifts, substantial heterogeneity across the individual forecasts may survive the tournament selection process, which results in more dispersed prices.

By contrast, by superposing Figures 3d and 1d, it appears that the welfare loss due to the dispersion of the individual consumption levels under SL – $\hat{\gamma}_{j,t}$ – does not move alongside the variance of aggregate output. In fact, a Taylor rule that emphasizes the stabilization of output instead generates a substantial amount of dispersion in the individual consumption levels. To understand why, it is necessary to go back to the microfoundations of the model with SL expectations. Following Arifovic et al. [forthcoming], Grimaud et al. [2024], we rewrite the consumption of agent j as

$$\widehat{c}_{j,t} = \mathbb{E}^*_{j,t}(\widehat{y}_{t+1}) - \sigma^{-1}(\hat{\imath}_t - \mathbb{E}^*_{j,t}(\widehat{\pi}_{t+1})) + \widehat{g}_t. \tag{41}$$

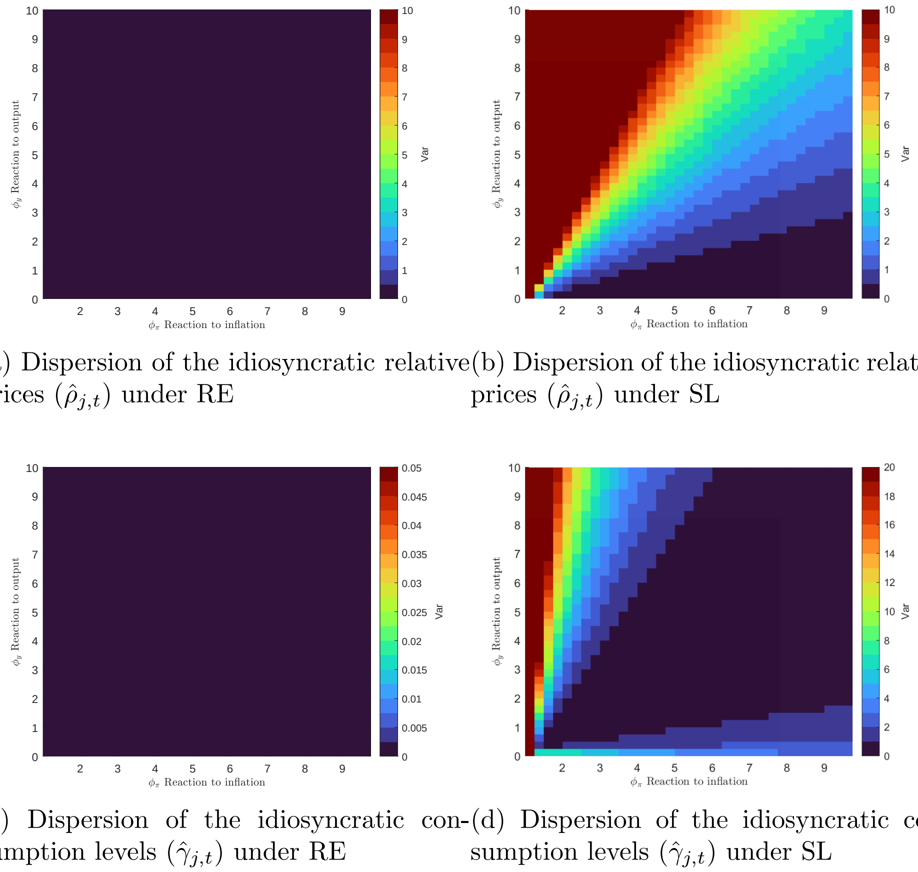

(a) Dispersion of the idiosyncratic relative prices ($\hat{\rho}_{j,t}$) under RE

(b) Dispersion of the idiosyncratic relative prices ($\hat{\rho}_{j,t}$) under SL

(c) Dispersion of the idiosyncratic consumption levels ($\hat{\gamma}_{j,t}$) under RE

(d) Dispersion of the idiosyncratic consumption levels ($\hat{\gamma}_{j,t}$) under SL

Notes: See Fig. 1.

Figure 3: Monetary policy trade-off with respect to the idiosyncratic components of welfare

Similarly, their idiosyncratic inflation rate is equivalent to:

$$\widehat{\pi}_{j,t} = \beta \mathbb{E}^*_{j,t}(\widehat{\pi}_{t+1}) + \kappa \widehat{y}_t + \widehat{u}_t. \tag{42}$$

From Eq. (42), we can immediately see that idiosyncratic inflation rates only depend on idiosyncratic inflation expectations (besides aggregate output and cost-push shocks). Hence, more heterogeneous inflation expectations entail more dispersed prices and stabilizing inflation leads to more homogeneous expected and, consequently, realized idiosyncratic inflation rates. However, Eq. (41) indicates that idiosyncratic consumption levels are not only determined by individual output expectations but also by individual inflation expectations. Therefore, two opposite effects are at play in determining consumption dispersion. On the one hand, in the same way, as volatile inflation generates more dispersed inflation expectations and prices, more volatile output leads to more dispersion in individual consumption levels. Yet, on the other hand, more volatile inflation – and its corollary of more heterogeneous inflation expectation – also increases the dispersion among agents' consumption levels. Therefore, either way, whether the CB puts a strong emphasis on output stabilization (which amounts to choosing a high ϕ_y value and a low ϕ_π value) or on inflation stabilization (with a high ϕ_π and a low ϕ_y), the dispersion of individual consumption levels goes up (to see that, look at these two warm-colored areas in Fig. 3d). Consequently, the welfare loss arising from the dispersion of idiosyncratic consumption levels is minimized for intermediary cases where the reaction coefficients on inflation and output gaps are of similar magnitude; see the darker areas in the parameter space in Fig. 3d that tend to gather around the first diagonal.

Before turning to a complete welfare analysis, it is worth contrasting our framework with the standard adaptive learning models. Heterogeneous expectations under SL exacerbate the welfare loss with respect to the RE benchmark *via* two channels. The first channel is the excess volatility caused by the persistence in the aggregate variables (inflation and output gaps here) introduced by the learning process, while under RE and a standard contemporaneous IT rule, endogenous variables are jump variables. This first channel is also present in models of adaptive learning or under backward-looking expectations in general. The second channel through which SL induces welfare losses, however, is a property of heterogeneous expectations in our micro-founded model. It stems from the costs of dispersion in prices and consumption levels that result from heterogeneous expectations. Therefore, heterogeneous expectations complicate the stabilization trade-off of the CB because a higher aggregate volatility turns into more heterogeneous individual expectations which results in more dispersion in firms' prices and households' consumption choices. Furthermore, the dispersion in consumption levels increases with the dispersion in both inflation and output expectations, so that no matter which side of the trade-off the CB chooses to favor (output or inflation), consumption levels are hard to align.

4.3.3 Welfare implications

Putting together the insights from the aggregate and the idiosyncratic components into the welfare criterion, Fig. 4 shows the welfare outcomes under IT ($\omega = 1$) in the parameter space $\{\phi_\pi, \phi_y\} = \{(1, 10], (0, 10]\}$ under RE (Fig. 4a) and SL (Fig. 4b). Unsurprisingly, due to the absence of dispersion in consumption and prices, the welfare under RE is strikingly higher than under SL – note that the two sub-figures are on different scales, the scale under SL is much wider than under RE. The welfare differences between the two expectation schemes are mostly in terms of level of welfare: in both cases, the high welfare zone corresponds to a strong reaction to inflation and a more moderate reaction to the output gap. We should also note that the parameter space where the welfare is maximized under SL is larger than under RE, although a lower welfare is obtained than under RE overall (compare the size of the dark-red triangle-shaped areas in Fig. 4a and 4b).

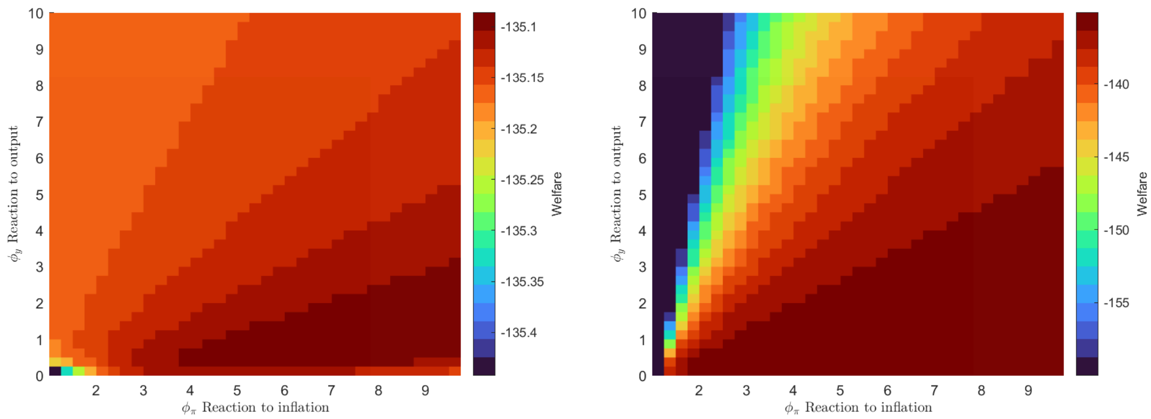

(a) Welfare under RE

(b) Welfare under SL

Notes: See Fig. 1. Note the difference in scale between the two sub-figures.

Figure 4: Welfare under IT

4.4 Optimal policy rules results

Over the recent years, history-dependent policy rules have gained interest. For instance, in August 2020, the Fed announced that the inflation target would be 'an average over time.' To account for these refinements, in what follows we consider optimal Taylor rules over the full set of policy parameter values. Specifically, we use $\{\phi_\pi, \phi_y, \omega\} = \{(1, 10], (0, 10], [0, 1]\}$. In the toolbox, the function `OSR_sl` finds these parameter values that maximize $\mathcal{W}_t$.[26]

Under SL, an additional layer of complexity with respect to the optimization problem under RE stems from the idiosyncratic random draws and the non-linear feedback loops between the SL expectations and the realization of the endogenous variables. For this reason, it is advisable to use a global optimization algorithm.[27] Another hurdle to this optimization problem is that in the absence of financial or investment frictions, the variance of the interest rate does not affect the welfare of the agent. Woodford [2011] notes that this has unpleasant implications for the design of optimal monetary policy rules. Indeed, in a simple three-equation models, the welfare-maximizing strategy consists of a very large and totally implausible set of coefficients $\{\phi_\pi, \phi_y\}$. In order to avoid this results, we follow the spirit of Woodford [2011] and an ad-hoc penalty $-\gamma_i Var(\hat{\imath}_t)$ - with $\gamma_i = 9$ - to the welfare function 39. Thus, when optimizing, the CB balances the marginal gain in aggregate volatility with the marginal loss in financial volatility. The calibrated value is chosen such that the $\phi_\pi \approx 1.5$ when optimized under RE. Again, the point of the exercise is not to generate normative implications for policymakers but to contrast the policy prescriptions under RE and SL in a calibrated small-scale model.

[26] The function can also be set to minimize $\mathcal{L}(\alpha)$.

[27] The toolbox uses an updated version of the `Fminsearch` function of MATLAB [Inc., 2022] – which is based on the Nelder-Mead simplex minimization algorithm – that accepts constraints on parameters in

Scenario: ***Optimal values*** $\{\phi_\pi, \phi_y, \omega\}$	(I) ***Benchmark IT*** $\{1.5, 0.125, 1\}$	(II) ***Optimal RE rule*** $\{1.53, 0.5, 0\}$	(III) ***Optimal SL rule*** $\{9.42, 3.25, 0\}$
Inflation variance $var(\hat{\pi}_t)$:			
RE model	0.2580	0.0709	0.0661
SL model	0.9311	0.0989	0.0913
Output gap variance $var(\hat{\pi}_{j,t})$:			
RE model	1.637	2.3168	2.32
SL model	2.164	4.049	3.7192
Variance of individual consumption levels $var(\hat{\gamma}_{j,t})$:			
RE model	0	0	0
SL model	1.5090	0.5449	0.3747
Variance of individual price dispersion $var(\hat{\rho}_{j,t})$:			
RE model	0	0	0
SL model	1.3457	0.0693	0.0628
Welfare $(\mathcal{W}_t)$:			
RE model	-135.25	-135.04	-135.05
SL model	-136.49	-135.16	-135.14
Loss in permanent consumption equivalent w.r.t the deterministic steady state			
RE model	0.1011%	0.0373%	0.0388%
SL model	0.4788%	0.0732%	0.0677%

Notes: See Fig. 1.

Table 1: Simulated moments under different CB's reaction functions

Table 1 contrasts the simulated moments of the model under the standard IT policy previously considered in Section 4.3 in Col. (I), the optimal policy under RE in Col. (II), and the optimal policy under SL in Col. (III). It is worth recalling that the textbook RE model under IT (which corresponds to $\omega = 1$) does not contain any lagged state variables (such as those stemming from consumption habits, price indexation or additional real or nominal frictions in more sophisticated versions of the NK model). In this baseline version, only when persistence is introduced in the monetary policy rule *via* history-dependent inflation target (i.e. $\omega < 1$) – are inflation and output gaps no longer only jump variables. By contrast, under SL, even under a simple IT rule, expectations are dependent on past realizations of the endogenous variables through the computation of the fitness of each forecast and the tournament selection. Therefore, under SL, inflation and output gaps exhibit endogenous persistence no matter the specification of monetary policy.

Col. (I) of Table 1 is in line with the analysis of Section 4.3: all aggregate moments under a standard IT regime exhibit a larger variance with SL expectations than with RE. and consequently, welfare under SL is considerably smaller than under RE. Looking into optimal policy design in Cols. (II) and (III), we find the common result that PLT is

the optimization space; see functions `Fminsearchbnd`, `OSR_run`, and `OSR_opti` in the toolbox.

optimal under RE (Col. II). The corresponding set of policy parameters that maximize welfare is $\{\phi_\pi, \phi_y, \omega\} = \{1.53, 0.5, 0\}$, which refers to an aggressive PLT rule with respect to output.[28] The variance of inflation is greatly reduced whereas output volatility goes up. This is a standard result reflecting the fact that in these simple models, welfare losses are mostly driven by inflation volatility rather than by output volatility.

Interestingly, using the optimal policy rule under RE in a model with SL expectations does result in less volatile inflation than under the standard IT rule in Col. I, albeit still higher than under RE. By contrast, optimal monetary policy under RE results in higher output volatility than under a standard calibrated IT rule, and the increase is steeper under SL than under RE. This result reflects the trade-off between inflation and output gap stabilization illustrated in the efficiency frontiers 2 and the higher sacrifice ratio under SL than under RE. This higher output variance is accompanied by more dispersion in output gap expectations, which can lead to more dispersed consumption levels. However, a lower inflation variance also leads to less dispersed inflation expectations, which has the opposite effect on consumption dispersion (see, again, Eq. (41)). This second effect dominates, and consumption patterns are less heterogeneous under the optimal RE policy rule than under the calibrated IT rule. Taken together, these observations result in the optimal policy rule under RE being welfare-improving under both expectation schemes, not only under RE.

The optimal policy rule under SL is reported in Col. (III) of Table 1 and corresponds to a more aggressive price level targeting rule than under RE, with a strong reaction to output gap, namely $\{\phi_\pi, \phi_y, \rho, \omega\} = \{9, 42, 3.25, 0\}$. Under the optimal rule under SL, inflation is better stabilized and prices are less dispersed than under a standard IT regime, although not as well as under the optimal RE policy. The improvement with respect to the RE optimal policy rather comes from less dispersed consumption patterns under the SL optimal rule. It is noteworthy to remark that output stabilization is still worse than under the benchmark case but better than under the optimal RE rule. To conclude, under both expectation schemes, the SL optimal policy is welfare improving with respect to the standard calibrated IT regime.

Finally, let us briefly compare our results with the ones obtained under adaptive learning literature. Under adaptive learning, reaction coefficients also require higher values than under RE [Orphanides and Williams, 2005] and this seems to also be true under SL. However, history-dependent rules have been found to generate *more* volatility under backward-looking expectations such as adaptive learning than standard IT rules (see Honkapohja and McClung 2023, Amano et al. 2020). By contrast, under SL, our results seem to indicate that optimal policy rules are robust to this form of heterogeneous expectations: history-dependent policy rules are welfare-improving both under RE and SL.

[28] Recall that the value of ϕ_π is a construction from the calibration of γ_i, the loss term introduced to penalize large movements in interest rates.

This result is good news for policymakers who are uncertain about how agents form their expectations. Our results also provide a rationale for history-dependent rules even in the presence of learning and heterogeneity in expectations.

5 Conclusion

This paper aims to formulate a standard micro-founded general-equilibrium model with heterogeneous expectations that are shaped by idiosyncratic news and social interactions. We derive a general solution to this class of models that we implement in a Dynare toolbox to simulate macroeconomic models under SL expectations and contrast their properties with their RE counterparts. We illustrate such HENK models within the context of the textbook three-equation NK model and analyze optimal history-dependent monetary policy rules under SL and RE. We find that history dependence in the inflation target is also optimal under SL. Moreover, such policy rule is robust to the expectation specification in the sense that both optimal policies under SL and under RE are welfare-improving no matter whether expectations are rational or heterogeneous.

While this class of expectation models based on the functioning of GAs and evolutionary learning has been used for a long time, a general solution method and a toolbox both constitute substantial improvements upon the existing literature. We hope that our present contribution will help facilitate the use of these models among scholars and practitioners alike by formulating these models in standard recursive expressions, deploying them in the common `Dynare` framework, and considerably diminishing the computation burden of these simulation models. The toolbox may also facilitate the replication of the results.

Our contribution opens up many research avenues. In particular, SL expectations are a promising mechanism to simulate so-called animal spirits or belief contagions in markets in larger-scale macroeconomic models. The toolbox further opens up the possibility of estimating HENK models with full-information methods.

References

Klaus Adam and Albert Marcet. Internal rationality, imperfect market knowledge and asset prices. Journal of Economic Theory, 146(3):1224–1252, 2011. 3, 7

Robert A Amano, Stefano Gnocchi, Sylvain Leduc, and Joel Wagner. Average is good enough: Average-inflation targeting and the ELB. Technical Report 2020-31, Bank of Canada Staff Working Paper, 2020. 29

Philippe Andrade, Gaetano Gaballo, Eric Mengus, and Benoit Mojon. Forward guidance and heterogeneous beliefs. American Economic Journal: Macroeconomics, 11(3):1–29, 2019. 3

Jasmina Arifovic. Evolutionary algorithms in macroeconomic models. Macroeconomic Dynamics, 4(3):373–414, 2000. 1

Jasmina Arifovic, James Bullard, and Olena Kostyshyna. Social learning and monetary policy rules. Economic Journal, 123(2001):38–76, 2013. 1, 2, 4, 5, 6, 7, 8, 14

Jasmina Arifovic, Stephanie Schmitt-Grohé, and Martın Uribe. Learning to live in a liquidity trap. Journal of Economic Dynamics and Control, 89:120–136, 2018. 1, 21

Jasmina Arifovic, Alex Grimaud, Isabelle Salle, and Gauthier Vermandel. Social Learning and Monetary Policy at the Effective Lower Bound. Journal of Money Credit and Banking, forthcoming. 1, 9, 19, 21, 24

Florin O Bilbiie. The new keynesian cross. Journal of Monetary Economics, 114:90–108, 2020. 3

William A. Branch and Bruce McGough. Dynamic predictor selection in a new Keynesian model with heterogeneous expectations. Journal of Economic Dynamics and Control, 34(8):1492–1508, 2010. 3

Flora Budianto, Taisuke Nakata, and Sebastian Schmidt. Average inflation targeting and the interest rate lower bound. European Economic Review, 152:104384, 2023. 20

Carlos Carvalho, Stefano Eusepi, Emanuel Moench, and Bruce Preston. Anchored inflation expectations. American Economic Journal: Macroeconomics, 15(1):1–47, 2023. 3

Olivier Coibion and Yuriy Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review, 105(8):2644–2678, 2015. 1

Stefano Eusepi and Bruce Preston. Expectations, learning, and business cycle fluctuations. American Economic Review, 101(6):2844–2872, 2011. 3

Stefano Eusepi and Bruce Preston. Fiscal foundations of inflation: imperfect knowledge. American Economic Review, 108(9):2551–2589, 2018a. 4

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

George Evans. Expectational stability and the multiple equilibria problem in linear rational expectations models. Quarterly Journal of Economics, 100(4):1217–1233, 1985. 2

George W Evans and Seppo Honkapohja. Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models. Econometrica, pages 195–206, 1995. 2

George W. Evans and Seppo Honkapohja. Learning and Expectations in Macroeconomics. Princeton University Press, 2001. ISBN 9780691049212. 2

George W. Evans, Eran Guse, and Seppo Honkapohja. Liquidity traps, learning and stagnation. European Economic Review, 52(8):1438–1463, 2008. ISSN 00142921. doi: 10.1016/j.euroecorev.2008.05.003. 3

Alex Grimaud, Isabelle Salle, and Gauthier Vermandel. A Note on the Microfoundations for the Three-Equation New Keynesian Model under Social Learning. 2024. mimeo. 19, 20, 24

J. H. Holland. Adaptation in Natural and Artificial Systems. Ann arbor: University of michigan press edition, 1975. 1

Cars Hommes. Behavioral and experimental macroeconomics and policy analysis: A complex systems approach. Journal of Economic Literature, 59(1):149–219, 2021. 1, 2

Seppo Honkapohja and Nigel McClung. On robustness of average inflation targeting. Available at SSRN 4021712, 2023. 29

The MathWorks Inc. Matlab version: 9.13.0 (r2022b), 2022. URL https://www.mathworks.com. 18, 27

Michel Juillard et al. Dynare: A program for the resolution and simulation of dynamic models with forward variables through the use of a relaxation algorithm, volume 9602. Citeseer, 1996. 2, 18

Greg Kaplan, Benjamin Moll, and Giovanni L Violante. Monetary policy according to hank. American Economic Review, 108(3):697–743, 2018. 3

Sharon Kozicki and Peter A Tinsley. Effective use of survey information in estimating the evolution of expected inflation. Journal of Money, Credit and Banking, 44(1):145–169, 2012. 4

Guido Lorenzoni. Optimal monetary policy with uncertain fundamentals and dispersed information. Review of Economic Studies, 77(1):305–338, 2010. 13

Albert Marcet and Thomas J Sargent. Convergence of least squares learning mechanisms in self-referential linear stochastic models. Journal of Economic theory, 48(2):337–368, 1989. 2

Domenico Massaro. Heterogeneous expectations in monetary DSGE models. Journal of Economic Dynamics and Control, 37(3):680–692, 2013. ISSN 01651889. 3

Roland Meeks and Francesca Monti. Heterogeneous beliefs and the phillips curve. Journal of Monetary Economics, 2023. 1

Athanasios Orphanides and John C. Williams. Inflation scares and forecast-based monetary policy. Review of Economic Dynamics, 8(2):498–527, 2005. 24, 29

Pablo Ottonello and Thomas Winberry. Financial heterogeneity and the investment channel of monetary policy. Econometrica, 88(6):2473–2502, 2020. 3

John B Taylor. Discretion versus policy rules in practice. 39:195–214, 1993. 20

Harald Uhlig. A toolkit for analyzing nonlinear dynamic stochastic models easily. 1995. 34

Michael Weber, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion. The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications. Journal of Economic Perspectives, 36(3):157–84, 2022. 1

Michael Woodford. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press, 2011. 27

A The derivation of the general solution method

Following Uhlig [1995] and using our notation (with $D_1 = I$ and $D_2 = P + I$ in Eq. (27)), the model is as follows

$$F\mathbb{E}_t z_{t+1} + FD_2\varphi_t + Gz_t + Hz_{t-1} + M\varepsilon_t = 0;$$

and by definition, we assume the solution follows this state-space form

$$z_t = Pz_{t-1} + Q\varepsilon_t + R\varphi_t.$$

The model consistent forecast can be written in this fashion

$$\begin{aligned}
\mathbb{E}_t z_{t+1} &= P\left(Pz_{t-1} + Q\varepsilon_t + R\varphi_t\right) + Q\mathbb{E}_t\varepsilon_{t+1} + R\mathbb{E}_t\varphi_{t+1}, \\
&= P^2 z_{t-1} + PQ\varepsilon_t + PR\varphi_t + 0 + R\varphi_t, \\
&= P^2 z_{t-1} + PQ\varepsilon_t + (P + I)R\varphi_t.
\end{aligned}$$

Plugging back the forecast into the ALM of the model

$$\begin{aligned}
0 &= F(P^2 z_{t-1} + PQ\varepsilon_t + (P + I)R\varphi_t) + FD_2\varphi_t + Gz_t + Hz_{t-1} + M\varepsilon_t, \\
-Gz_t &= (FP^2 + H)z_{t-1} + FPQ\varepsilon_t + F(P + I)R\varphi_t + FD_2\varphi_t + M\varepsilon_t, \\
-G(Pz_{t-1} + Q\varepsilon_t + R\varphi_t) &= (FP^2 + H)z_{-1} + FPQ\varepsilon_t + F\left(\left(P + I\right)R + D_2\right)\varphi_t + M\varepsilon_t.
\end{aligned}$$

Using the undetermined coefficients method we can assume that P solves for

$$\begin{aligned}
-GP = (FP^2 + H) &\Leftrightarrow -GP - FP^2 = H \\
&\Leftrightarrow -\left(FP + G\right)P = H \\
&\Leftrightarrow P = -\left(FP + G\right)^{-1} H
\end{aligned}$$

and then

$$\begin{aligned}
-GQ = FPQ + M &\Leftrightarrow -GQ - FPQ = M \\
&-\left(FP + G\right)Q = M \\
&Q = -\left(FP + G\right)^{-1} M
\end{aligned}$$

and then

$$
\begin{aligned}
-GR = F\left((P+I)\,R + D_2\right) &\Leftrightarrow -GR - FPR - FR = FD_2 \\
&\Leftrightarrow (-G - FP - F)R = FD_2 \\
&\Leftrightarrow R = (-G - FP - F)^{-1} FD_2 \\
&\Leftrightarrow R = (-G - FP - F)^{-1} F(P+I).
\end{aligned}
$$

B Representative run of the model

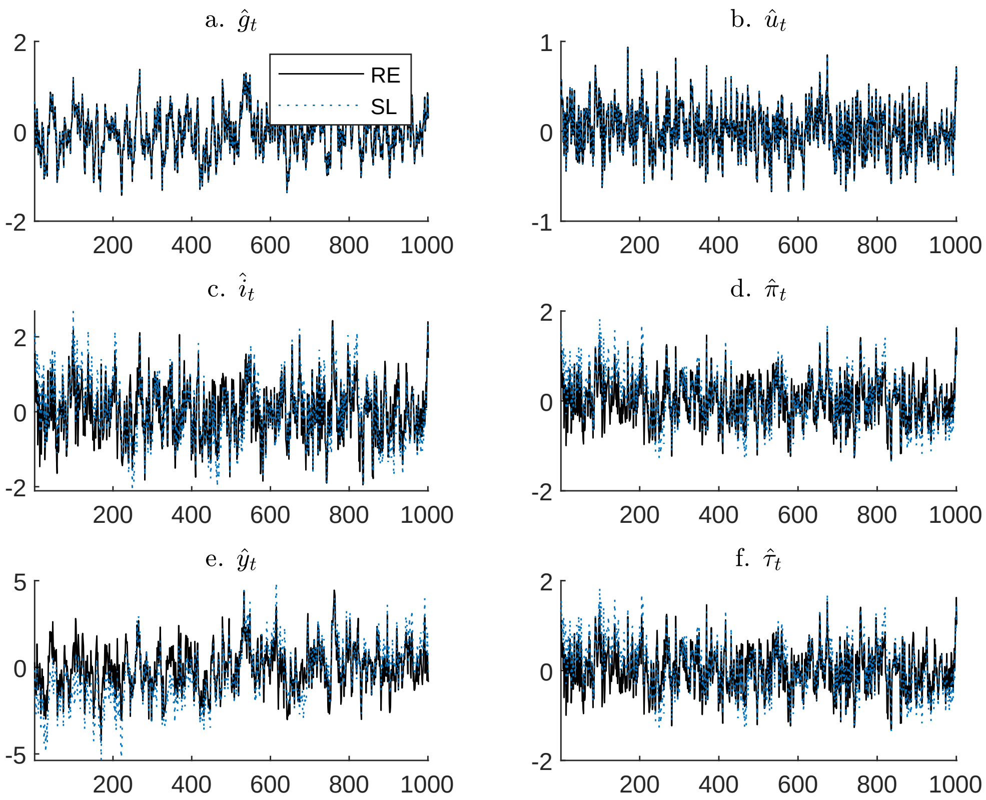

Notes: We removed the first 300 periods of initialization of the SL model.

Figure 5: Representative run of the model under SL and RE

C Response to a demand shock

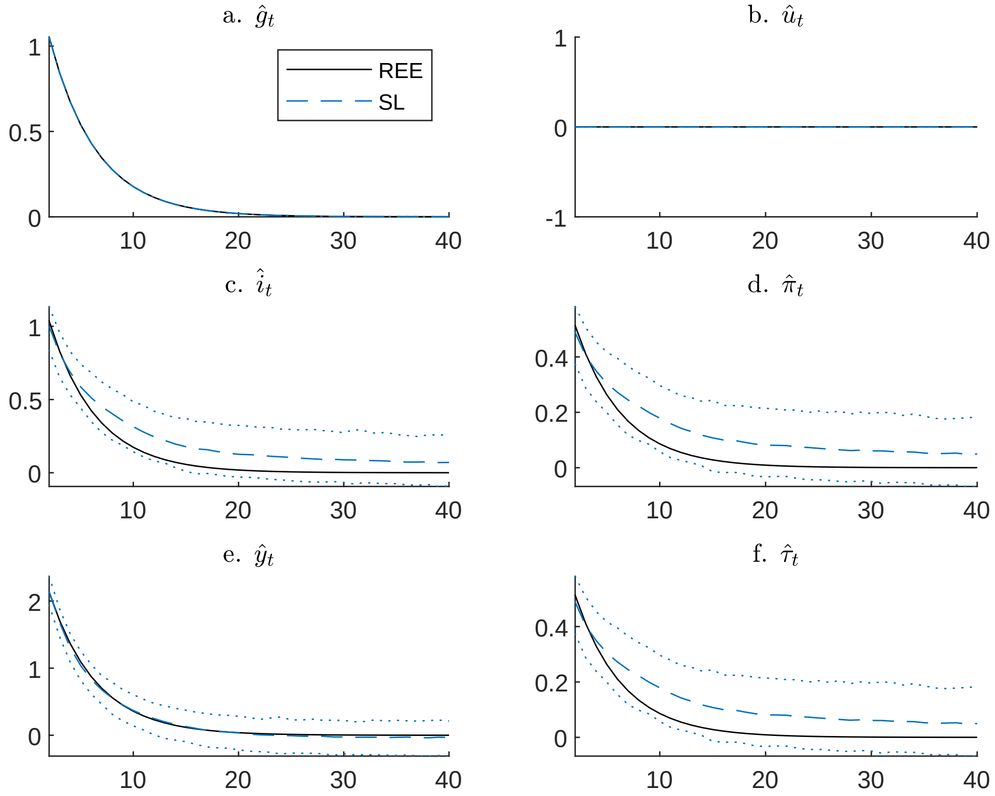

Notes: For SL, we plot the mean,5%, and 95% confidence intervals over 1000 SL Monte Carlo simulations with different initial histories for 300 periods and different idiosyncratic SL draws for the mutations and the tournaments over the full simulation. Note the additional persistence from the unanchoring of inflation and inflation expectations under SL versus under RE.

Figure 6: IRFs to a positive +1%, demand shock under RE and SL

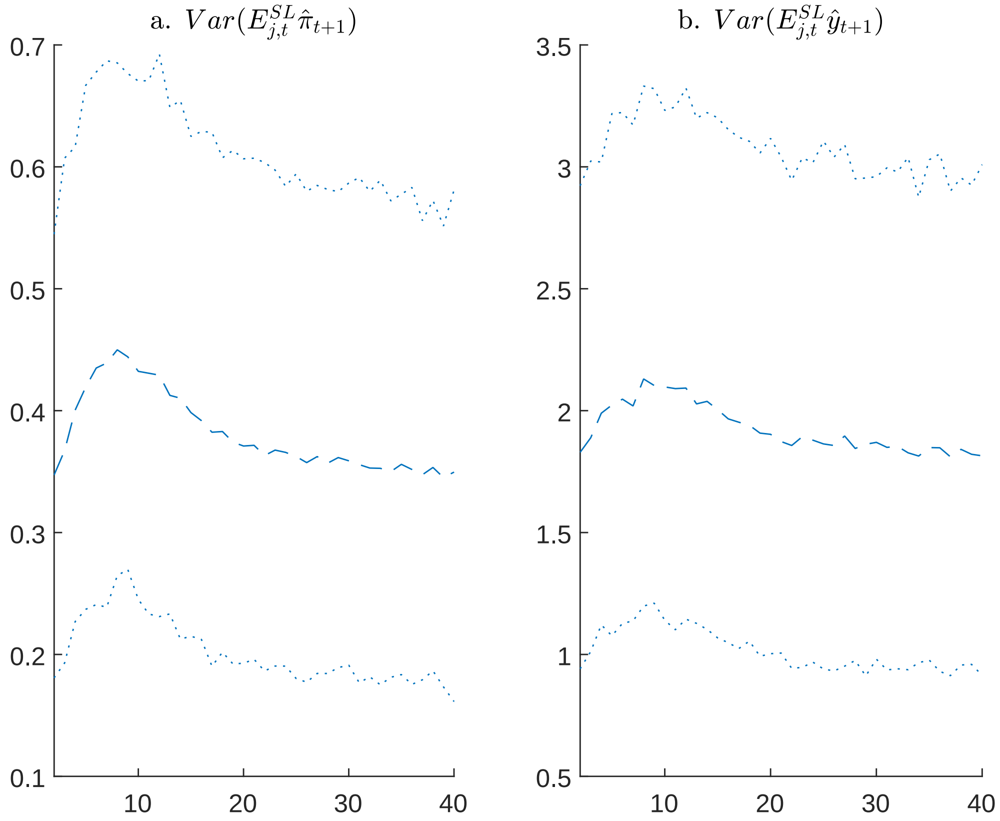

Notes: See Fig. 6. There is no RE counterpart because there is no cross-sectional dispersion with a representative agent. Note how the disagreement across agents increases in the wake of the shock.

Figure 7: Cross-sectional dispersion in individual SL expectations for the IRFs to a positive +1%, demand shock