Social learning expectations: microfoundations and a Dynare toolbox

Grimaud, Alex; Salle, Isabelle; Vermandel, Gauthier

DOI:
10.57938/70a0ca9f-2669-4934-b43c-a888b5ae1dac

Published: 01/07/2023

Document Version
Publisher's PDF, also known as Version of record

Department of Economics
Working Paper No. 339

Social learning expectations: microfoundations and a Dynare toolbox

Alex Grimaud
Isabelle Salle
Gauthier Vermandel

July 2023

Social learning expectations: microfoundations and a Dynare toolbox*

Alex Grimaud[†] Isabelle Salle[‡] Gauthier Vermandel[§]

May 3, 2023

Abstract

Social learning (SL) is a behavioral model in which expectations and the resulting aggregate dynamics stem from the interactions of a large amount of heterogeneous agents. Nonetheless, this framework has so far lacked micro-foundations and a general-solution method. This paper bridges these two gaps with: (i) a micro-founded New Keynesian model with social learning expectations; (ii) a general solution method that we implement in a Dynare toolbox that solves any linear state-space model with SL expectations. The resulting framework provides a self-contained tool to contrast policy analysis under SL and rational expectations. As an illustration, optimal monetary policy rules are studied under the two expectation regimes.

***Keywords*:** Inflation targeting, Monetary policy and Heterogeneous expectations.
***JEL codes*:** E32, E52, E58 and E71.

1 Introduction

A major contribution of late Prof. Jasmina Arifovic is the application of genetic algorithms to model decisions and expectations in economic models. Due to their nature, these algorithms represent heterogeneous and boundedly-rational expectations that result in non-linear learning dynamics. One fruitful application is the use of these algorithms to model the evolution of beliefs under social learning (SL). Throughout her career, this class of learning models and the linear rational expectations (RE) representative agent models used as a benchmark in macroeconomic and monetary analysis have remained separated by a thick border. This paper aims to address this discrepancy by applying

*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).

the tools of the benchmark framework to Jasmina's SL model. The result is a micro-founded heterogeneous-expectation New Keynesian (HENK) framework deployed in a `Dynare` toolbox for monetary policy analysis.

Expectations constitute a major transmission channel of macroeconomic policies to the economy because they play a key role in agents' economic and financial decision-making. In macroeconomic models, expectations have been usually modeled within the full-information rational-expectations (FIRE) framework. However, the FIRE benchmark abstracts from important features of expectation formation as highlighted by empirical evidence from laboratory or survey data. One of these features is the heterogeneity in expectations, i.e. a substantial cross-sectional dispersion in forecasts. Yet, heterogeneity has been found to be a pervasive feature of real-world expectations, whether these expectations are elicited from professional forecasters, households, firms, managers or policymakers themselves.[1] Disagreement between agents may arise for a variety of reasons; for instance, because agents have different cognitive abilities (in particular, different depths of reasoning or levels of sophistication), use different information sets (for example, if they pay a different level of attention to news or use different sources of information), or if agents are endowed with different models of the economy.

By neglecting this disagreement, the FIRE models ignore frictional coordination as a source of economic fluctuations. This paper develops a HENK framework that can account for this information friction, quantify its importance in terms of welfare loss with respect to the FIRE benchmark and analyze optimal monetary policy when heterogeneous expectations cause excess volatility in inflation and output.[2][3]

Our HENK framework relies on social learning (SL, hereafter) among agents. This class of learning models is based on the evolutionary principle of 'survival of the fittest' and is implemented using variations of genetic algorithms initially developed for optimization purposes [Holland, 1975]. In detail, these models combine two forces: i) the individual exploration of the space of strategies and ii) the collective exploitation of the existing strategies. An economic interpretation of this class of algorithms can be a setting where agents receive idiosyncratic news (or private signals) and update their forecasts by adopting the most accurate ones among their peers during social interactions. While models with

[1]See, *inter alia*, Coibion and Gorodnichenko [2015], Hommes [2021], Weber et al. [2022].

[2]See Arifovic et al. [2023] for a survey of the literature exploring the interplay between heterogeneity and genetic algorithm in macroeconomic models.

[3]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 in [see e.g Kaplan et al., 2018, Bilbiie, 2020, Ottonello and Winberry, 2020]. Moreover, in HENK models, beliefs' heterogeneity is dynamic and the distribution of the different beliefs varies 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 [see e.g Reiter, 2009]

SL expectations have been studied in the literature for a long time[4], these applications lack four inter-related features to integrate them within the toolkit of the mainstream macroeconomists, namely: i) a general-solution method to write any dynamic stochastic general equilibrium (DSGE) model under SL; ii) micro-foundations of the expectation formation process; iii) a welfare criterion to perform policy analysis and iv) a `Dynare` toolbox to simplify the implementation of the SL framework and the replication of the results. This paper briges these gaps and brings these four features.

In Section 2, we derive a general solution method for micro-founded macroeconomic models where expectations are the result of an SL process. Our method introduces a new level of consistency in the solution of SL models because agents' perceived laws of motion internalize the future changes of their idiosyncratic expectations. We then present an application with micro-founded SL expectations within the textbook three-equation New Keynesian (NK) model in Section 3 (the derivation of the welfare criterion is deferred to Appendix B and is taken from Arifovic et al. 2023).

In Section 4[5], we introduce a `Dynare` [Juillard et al., 1996] toolbox for Matlab to simulate standard macroeconomic models with SL expectations. The toolbox is flexible and can be utilized to simulate any standard dynamic and stochastic models under SL expectations. We provide general functions to generate simulate time series and IRFs. Within the context of the baseline NK model, the toolbox also allows one to perform welfare analysis under RE and under SL and for various monetary policy rules.

As an illustration, Section 5 contrasts optimal monetary policy under RE and under SL expectations. In particular, we show that the stabilization trade-off of the central bank (CB) is made more challenging under heterogeneous expectations than under RE because heterogeneous expectations create dispersion in prices and consumption levels that entails an additional welfare loss with respect to the RE benchmark with exogenous fluctuations only. We find that a monetary policy rule with a fair amount of history dependence and a rather aggressive coefficient on inflation developments is optimal under SL. Moreover, such policy rule is robust to the expectation scheme in the sense that no matter whether policies are optimized assuming RE or SL, they are welfare-improving under both expectation schemes.

2 A general formulation and resolution

This section provides a general formulation of a model under SL expectations and its general solution based on perturbation methods.

[4]See Arifovic [2000] for a survey of early contributions and Arifovic et al. [2013, 2018, 2023] for recent examples.

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

2.1 Definitions

To introduce the definition of SL expectations, let us first define an economy under the FIRE benchmark.

Definition 1. *An economy populated by J agents, indexed by $j = 1, ..., J$, endowed with full information and rational expectations is modeled as an environment in which all agents know the structure of the economy, share the same information in all periods, face no uncertainty about one another's beliefs and behavior (i.e. there is no strategic uncertainty), and reach a perfect consensus about the current state and the future prospects of the economy [Angeletos et al., 2018]. Formally, each agent j forms expectations about the future state of the economy $\mathbb{E}_{j,t}(y_{t+1}) \equiv \mathbb{E}_j(y_{t+1}|\Omega_t)$ conditional on the common information set Ω_t, which equalizes all agents' expectations on the following value:*

$$\mathbb{E}_j(y_{t+1}|\Omega_t) \equiv \mathbb{E}(y_{t+1}|\Omega_t), \forall j \in \{1, 2, ..., J\}, \quad (1)$$

where $\mathbb{E}_t$ is defined as the expectation operator after aggregating all agents' beliefs in our dynamic economy, which boils down to the rational expectation operator.

It is well known in macroeconomics that the FIRE benchmark (1) rules out possible disagreement between agents and, by doing so, ignore frictional coordination as a source of economic fluctuations. By contrast, the SL algorithm gives a central role to these information frictions in the evolution of endogenous variables. We define an economy with information frictions under SL expectations as follows:

Definition 2. *An economy populated by J agents, indexed by $j = 1, ..., J$, endowed with social learning expectations is modeled as an environment in which all agents act under bounded rationality, they know the structure of the economy, form heterogeneous beliefs about fundamentals based on their private signals, have limited information about one another's beliefs (i.e. there is strategic uncertainty), and fail to coordinate their expectations about the current state and the future prospects of the economy. Formally, each agent j forms expectations about the future state of the economy conditional on their idiosyncratic information set $\Omega_{i,t}$, which includes their private information $a_{j,t}$ about y_{t+1} and the past realizations of the endogenous variables, but exclude the private information of the $J-1$ other agents (i.e. $a_{k,t} \notin \Omega_{i,t}$, $k \neq j$). This information friction prevents the coordination of agents' expectations, which are then given by:*

$$\mathbb{E}_{j,t}(y_{t+1}) \equiv \mathbb{E}_j(y_{t+1}|a_{j,t}) \neq \mathbb{E}(y_{t+1}|a_{1,t}, a_{2,t}, ..., a_{J,t}) \text{ for } j \in \{1, 2, ..., J\}, \quad (2)$$

From Definition 2, one can note that the FIRE benchmark is nested in the SL environment: as soon as $a_{j,t} = E_t(a_{t+1})$, $\forall j, t$, i.e. if all agents have the same private information that further coincides with the rational expectations of the endogenous variables, all agents coordination on the REE. Under SL, the information friction stemming from the limited knowledge of each agent about one another's beliefs creates disagreement in forecasts. In the next subsection, we describe the formation process of the individual beliefs $a_{j,t}$ over time.

2.2 A general formulation of SL expectations

Consider an economy with a discrete population of agents, indexed by $j \in \{1, 2, ..., J\}$, who forecast the vector of forward-looking variables y_{t+1} of size $N_y \times 1$, where N_y is the number of endogenous variables. Each agent j has a private vector of information $a_{j,t}$ about y_{t+1} that is introduced in the model as:

$$\mathbb{E}_{j,t}\left(y_{t+1}\right) \equiv \mathbb{E}_{j}\left(y_{t+1}|\Omega_{j,t}\right) = \mathbb{E}_{j,t}\left(y_{t+1}\exp\left(a_{j,t}\right)\right). \tag{3}$$

The agent's specific information set is interpreted as some individual beliefs about the low-frequency component of y_{t+1}. For instance, in the NK models, the private signal $a_{j,t}$ about inflation can be interpreted as the relative anchorage of agent j's inflation expectation with respect to the CB's inflation target.

Under SL, individual beliefs are revised following a recursive process:

$$a_{j,t} - a_{j,t-1} = s\left(a_{j,t-1}, \xi\iota_{j,t}, y_{t-1}, \ldots, y_1\right) \text{ with } \iota_{j,t} \sim \mathcal{N}\left(0,1\right), \tag{4}$$

where ξ is a $N_y \times N_y$ covariance matrix of news, or "mutations" in the language of genetic algorithms, and $\iota_{j,t}$ some Gaussian disturbances. In this formulation, the function $s\left(\cdot\right)$ refers to the non-linear process of SL and is a three-step problem and encompasses two operators:

$$m_{j,t} = a_{j,t1} + \mathbb{1}_{z_{j,t}\leq\mu}\,\iota_{j,t}\xi \tag{5}$$

$$F_{j,t} = -\sum_{\tau=0}^{t}\rho^{\tau}\left(y_{t-1-\tau} - \mathbb{E}_{j,t-2-\tau}\left(y_{t-1-\tau}|m_{jt}\right)\right)^2$$

$$\left(a_{k,t}, a_{l,t}\right) = \mathbb{1}_{F_{k,t}>F_{k,t}}(m_{k,t}, m_{k,t}) + \mathbb{1}_{F_{k,t}\leq F_{l,t}}(m_{l,t}, m_{l,t}), \tag{6}$$

where Eq. (5) refers to the news process and Eq. (6) to the social interactions (or the tournament selection, in the language of genetic algorithms). In detail, $m_{j,t}$ denotes a new candidate (or *m*utation) in the forecasting rule of agent j; the frequency of mutations (or news arrival) in the population of agents, denoted by scalar μ, formally imposes that a number $J\mu$ of agents receive private news drawn from a Gaussian distribution $\iota_{j,t} \sim \mathcal{N}(0,1)$, where ξ is the covariance matrix of the idiosyncratic news shocks across the forward-looking variables. The subset of agents receiving news is selected based on a random variable $z_{j,t} \sim \mathcal{U}(0,1)$ drawn from a continuous uniform distribution. Each new vector candidate, $m_{j,t}$, is evaluated based on a discounted sum of forecast errors, denoted by $F_{j,t}$, computed as the squared difference between the previous observed realizations and the forecasts conditional on the information $m_{j,t}$ (see the second equation). Matrix ρ is a $N_y \times N_y$ diagonal matrix, with diagonal terms between 0 and 1. Note that $\mathbb{E}_{j,t-2}\left(y_{t-1}|m_{j,t}\right)$ is the forecast of y_{t-1} conditional on information set $m_{j,t}$ fed in the forecasting equation (3). Finally, Eq. (6) is the tournament, which pairs all agents into $J/2$ couples. The best forecasting coefficient within the couple is then adopted by the two members to form their expectations in period t.

Stacking all agents J into a vector $a_t = [a_{1,t}, a_{2,t}, ..., a_{J,t}]$, the individual beliefs of all agents follow a random walk:[6]

$$a_t - a_{t-1} = s\left(a_{t-1}, \xi\iota_t, y_{t-1}, \ldots, y_1\right) \text{ with } \iota_t \sim \mathcal{N}\left(0,1\right), \tag{7}$$

where ι_t is the $N_y \times J$ matrix of Gaussian mutations in period t.

Finally, the aggregation that maps idiosyncratic shocks to macro variables is given by:

$$\mathbb{E}_t\left(y_{t+1}|a_t\right) = \mathbb{E}_t\left(g\left(y_{t+1}, a_t\right)\right), \tag{8}$$

where $g\left(y_{t+1}, a_t\right)$ typically refers to the aggregation procedure. Under SL, the *every-voice-counts*-principle applies: $g\left(y_{t+1}, a_t\right) = 1/J \sum_{j=1}^{N} \mathbb{E}_{j,t}\left(y_{t+1}|a_{j,t}\right).$

Let us know embed the SL expectations into a general formulation of a DSGE model.

2.3 A general solution with perturbation methods

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

$$\mathbb{E}_t\left(f_\Theta\left(y_{t+1}, y_t, y_{t-1}, \varepsilon_t, a_t\right)\right) = 0, \tag{9}$$

where y_t is the vector of $N_y \times 1$ endogenous variables, ε_t is the vector of $N_\varepsilon \times 1$ stochastic innovations $\varepsilon_t \sim \mathcal{N}\left(0, \Sigma_\varepsilon\right)$ (i.e. the aggregate shocks) while a_t is a matrix $N_y \times J$ of random walks (i.e. the idiosyncratic shocks) with J the number of agents, f_θ is the vector of $N_y \times 1$ non-linear equations with structural parameters stacked into vector Θ. The term $\mathbb{E}_t\left(\cdot\right)$ denotes how the forward-looking variables are calculated based on the information set available in t.

In our setup, the idiosyncratic shocks affect the forward-looking variables only, which allows us to write the expected vector as a function $g(y_{t+1}, a_t)$ consistent with Equation

[6] Note that in what follows, we interpret $a_{j,t}$ as idiosyncratic shocks, but an alternative interpretation could be state variables.

(8). Our simplified system reads as:

$$\mathbb{E}_t \left\{ f_\Theta \left(g(y_{t+1}, a_t), y_t, y_{t-1}, \varepsilon_t \right) \right\} = 0. \tag{10}$$

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

$$\mathbb{E}_t \left(F \left(\hat{y}_{t+1} + I_N a_t U \right) + G \hat{y}_t + H \hat{y}_{t-1} + M \varepsilon_t \right) = 0, \tag{11}$$

where F, G and H are the $N_y \times N_y$ Jacobian matrices of $f_\Theta(\cdot)$ in y_{t+1}, y_t and y_{t-1} and M is the $N_y \times N_\varepsilon$ Jacobian matrix of system f_Θ in the vector of stochastic innovations ε_t . In this expression, $\hat{y}_t = y_t^{(1)} - \bar{y}$ is the first-order Taylor expansion term around the fixed point $\bar{y}$. As our expansion includes only first-order terms, $y_t^{(1)}$ is the linear-implied path consistent with the expansion order considered. Finally, $a_t U$ denotes the mean-wise aggregation procedure of the heterogeneous beliefs, with U a vector of size $1 \times J$, full of all $1/J$, resulting from the linearization of $g(\cdot)$.

The general solution of the dynamic linear system of equations with forward and backward-looking variables has the following form:

$$\hat{y}_t = P \hat{y}_{t-1} + Q \varepsilon_t + R a_t U, \tag{12}$$

with P, Q and R are three unknown matrices.

Iterating one step ahead, the expected aggregate and idiosyncratic shocks read as follows:

$$\mathbb{E}_t \left(\varepsilon_{t+1} \right) = 0, \tag{13}$$

$$\mathbb{E}_t \left(a_{t+1} \right) = a_t. \tag{14}$$

The first equation is consistent with the zero mean distribution of aggregate shocks. The second one assumes that (i) mutations are also on average zero (consistently with their distributions), (ii) the expected tournament outcomes are on average zero at an aggregate level. This assumption requires further discussion.

Given the law of motion in Eq. (7), what needs to be discussed in the aggregate forecast of the update of SL expectations, namely $\mathbb{E}_t \left(s \left(a_t, 0, \cdot \right) \right)$. In the absence of strategic uncertainty, each agent would anticipate the tournament output from $J(J-1)/2$ combinations in order to compute the change in the aggregate beliefs $s \left(a_t, 0, \cdot \right)$. By contrast, in the SL economy defined in 2, each agent forecasts under strategic uncertainty and has access to decentralized information, i. e. they ignore the beliefs of other agents outside their tournament pair. This decentralized and limited information does not allow agents to calculate the average forecast that would emerge from all possible pairs of all agents. Given the random-walk process of idiosyncratic a_j, agents naturally base their future forecasts conditional on the available private information they have as follows: $\mathbb{E}_t(a_{j,t+1}) = a_{j,t}$, which assumes, in turn, $\mathbb{E}_{j,t}(s(a_{jt}, 0, \cdot)) \simeq 0$. Stacking all the individual forecasts in a matrix yields to Eq. (14).[7]

As a consequence, at the idiosyncratic level and in the log-linearized model, beliefs may be written by

$$\mathbb{E}_{j,t}(\hat{y}_t) = P\hat{y}_{t-1} + Ra_{j,t} + Q\varepsilon_t, \quad \forall j, t. \tag{15}$$

Iterating forward, and using (13) and (14), it is possible to get

$$\begin{aligned} \mathbb{E}_{j,t}(\hat{y}_{t+1}) &= P\mathbb{E}_{j,t}(\hat{y}_t) + R\mathbb{E}_{j,t}(a_{j,t+1}) + Q\mathbb{E}_{j,t}(\varepsilon_{t+1}), \\ &= P\left(P\hat{y}_{t-1} + Ra_{j,t} + Q\varepsilon_t\right) + R\mathbb{E}_{j,t}(a_{j,t+1}) + Q\mathbb{E}_{j,t}(\varepsilon_{t+1}), \\ &= P\left(P\hat{y}_{t-1} + Ra_{j,t} + Q\varepsilon_t\right) + Ra_{j,t}, \\ &= P^2\hat{y}_{t-1} + (P + I)Ra_{j,t} + PQ\varepsilon_t. \quad \forall j, t. \end{aligned} \tag{16}$$

Aggregating Eq. (16) across all agents, the forward vector can be written as a linear projection on the current state and the idiosyncratic shocks as follow[8]

$$\begin{aligned} \mathbb{E}_t(\hat{y}_{t+1}) &= \frac{1}{J}\sum_{j=1}^{J}\mathbb{E}_{j,t}(\hat{y}_{t+1}) \\ &= P\hat{y}_t + R\frac{1}{J}\sum_{j=1}^{J}a_{j,t} \\ &= P\hat{y}_t + Ra_tU. \end{aligned} \tag{17}$$

It is critical to discuss the implications of Eqs. (16) and (17). Using the language of the adaptive learning literature, the perceived law of motion (PLM) of the agents is well-specified, i.e. it includes the same variables as the actual law of motion (ALM) of the economy given in Eq. (12), namely the past realizations of the endogenous variables y_{t-1} and the aggregate shocks ε_t. The PLM of the agents only differs with respect to the ALM of the economy when it comes to the intercept, which represents the steady-state values of the variables (zero in a model in gaps), in line with the use of steady-state learning. While the ALM depends on the aggregate value a_t, SL agents do not have access to the information of all the other agents and, therefore, only use their private information $a_{i,t}$ in their PLM. Nonetheless, agents internalize the evolution of theirs beliefs over time and recognize the influence of future SL dynamics over future realizations of the endogenous

[7] Note also that by the law of large number on J and because mutations are zero-mean, the limit case on J allows to pose $\lim_{J\to\infty}\mathbb{E}_t s(a_t, \xi\iota_{t+1}, \cdot) = 0$. This limit case highlights that the information set for all other households cannot be predicted and thus manipulated by one agent.

[8] We can then plug in for Eq. (12) in order to avoid expressing the model as a function of current state variables.

variables. The structure of the PLMs under SL is reminiscent of the concept of internal rationality developed in Adam and Marcet [2011], where expectations are a function of the expected changes in the PLM.

Solving the model (11) requires to find matrices P, Q and R assuming that the common component of expectations is as in Eq. (17). A policy function solving (11) must satisfy these three conditions:

$$\begin{aligned} P &= -\left(FP+G\right)^{-1}H, \\ Q &= -\left(FP+G\right)^{-1}M, \\ R &= -\left(FP+G\right)^{-1}F\left(R+I_N\right). \end{aligned} \tag{18}$$

Two of these matrices, namely P and Q, are exactly the same as in the REE. The only departure from the RE benchmark concerns R. Because P admits two solutions, there are also two corresponding matrices Q and 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 determinate for the toolbox to simulate the SL counterpart.

Unlike the existing designs of SL in the literature [see e.g Arifovic et al., 2013, 2018] that postulate a transition equation across $(t+1,t)$ different from the one across $(t,t-1)$, our solution method alleviates this issue and introduces a transition function that is consistent across time. Hence, the present solution enforces the law of iterated expectations because iterations of Eq. (17) remain valid across any forecasting horizon:

$$\mathbb{E}_t\left(\mathbb{E}_{t+1}\left(\hat{y}_{t+2}\right)\right) = \mathbb{E}_t\left(P^2\hat{y}_t + \left(P+I_{N_y}\right)Ra_tU\right) = \mathbb{E}_t\left(\hat{y}_{t+2}\right).$$

After a forward iteration of the policy rule in Eq. (16), one can rewrite the idiosyncratic forecast of an agent j as well as the fitness function as follows:

$$\begin{aligned} \mathbb{E}_{j,t}\left(\hat{y}_{t+1}\right) &= P\hat{y}_t + Ra_{j,t} \\ F_{j,t} &= -\sum_{\tau=0}^{t}\rho^{\tau}(Q\varepsilon_{t-1-\tau} + Ra_{t-1-\tau}U - Rm_{j,t})^2. \end{aligned}$$

Therefore, a good forecaster is a forecast $m_{j,t}$ that is able to minimizes predictions errors stemming from both the aggregate shocks ε and the idiosyncratic shocks a.

2.4 Determinacy under RE and stability under SL

Matrices P, Q, and R in Eqs (18) do not depend on the expectation scheme, they are identical under RE and SL.[9] Consequently, the determinacy conditions under RE – namely

[9] One way to think of matrix R under FIRE is to envision an SL model with $a_{j,t} = 0\ \forall j,t$. In the absence, of tournaments and thus learning behavior the model is surprisingly similar to dispersed

the parameter restrictions that ensure that the eigenvalues of the P matrix lie within the unit circle, i.e. $\max|\lambda_P| \in \mathbb{C} < 1$ – need to be satisfied to use the model under SL.[10]

To investigate stability under SL, one first need to rewrite the full dynamics of the model in a recursive way. Recall that the evolution of the idiosyncratic beliefs reads as:

$$a_t = a_{t-1} + s\left(a_{t-1}, \iota_t, \hat{y}_{t-1}, ..., \hat{y}_0\right).$$

One can rewrite $A_t \equiv a_t U$ in aggregate terms as follows:

$$\begin{aligned} A_t &= I_N A_{t-1} + s\left(a_{t-1}, \iota_t, \hat{y}_{t-1}\right) U \\ &= I_{N_y} A_{t-1} + S_t, \end{aligned}$$

where S_t is the compact update step that depends on both the aggregation of idiosyncratic forecasts and the selected news following the tournament.

One can then rewrite the policy rule in a recursive way:

$$\hat{y}_t = P\hat{y}_{t-1} + Q\varepsilon_t + RA_{t-1} + RS_t.$$

The full dynamic reads as:

$$\begin{bmatrix} A_t \\ \hat{y}_t \end{bmatrix} = \begin{bmatrix} I_{N_y} & 0_{N_y} \\ R & P \end{bmatrix} \begin{bmatrix} A_{t-1} \\ \hat{y}_{t-1} \end{bmatrix} + \begin{bmatrix} 1_{N_y} & 0_{N_y} \\ R & Q \end{bmatrix} \begin{bmatrix} S_t \\ \varepsilon_t \end{bmatrix},$$

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 (2.4) is stable under SL iif all eigenvalues of matrix $\Omega \equiv \begin{bmatrix} I_{N_y} & 0_{N_y} \\ R & P \end{bmatrix}$ lie within the unit circle. Matrix Ω is a triangular block matrix. Hence, we have $det(\Omega) = det(I_{N_y})det(P)$, or $det(\Omega - \lambda I_{N_y}) = det(1 - \lambda I_{N_y})det(P - \lambda I_{N_y})$, where the roots $\lambda \in \mathbb{C}$ are the eigenvalues of Ω. It follows that the eigenvalues of Ω are 1 (with multiplicity N_y) and the eigenvalues of P. Hence, the eigenvalues of P need to lie within the unit circle (which is equivalent to saying that the model (12) is determinate) for the model to be stable under SL. However, 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. □

Stability of the REE under SL may or may not result if the model is determinate and

information models à la Lorenzoni [2010].

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

depends on the non-linearities introduced by the SL expectations via the function $s(\cdot)$, for which no analytical solution is available.

3 The three-equation New Keynesian model with SL expectations

3.1 A micro-founded heterogeneous-expectation model

We develop here the textbook three-equation NK model with heterogeneous expectations shaped by social interactions. The time and the number of agents are discrete. Specifically, we assume that the economy is populated by an infinitely-living family composed by a discrete number J of members, indexed by $j \in [1, ..., J]$. Each member operates an intermediary-good-sector firm, so that indexes J indifferently refer to either firms or households, that we may then call agents, and firms and households have the same discount factor. These agents are identical (in particular in terms of preferences and technology, etc.) except when it comes to their inflation and output expectations that are heterogeneous.

3.1.1 Main assumptions on heterogeneity

Under the SL algorithm of Arifovic et al. [2013], the information set of the agents differs from the one prevailing under RE. Let x denote the variable that agents in the model need to forecast. Agent j's one-step-ahead expectation of x is given by:

$$\mathbb{E}_{j,t}\left\{x_{t+1}\right\} = \mathbb{E}_t\left\{x_{t+1}|a^x_{j,t}\right\} = \mathbb{E}_t\left\{\exp(a^x_{j,t})x_{t+1}\right\}, \tag{19}$$

where $\mathbb{E}_t\left\{x_{t+1}\right\}$ corresponds to the model-consistent expectation of variable x for $t+1$ given the information set available at time t – in other words, the rational expectation of x_{t+1} – and $\exp(a^x_{j,t})$ is the idiosyncratic information of agent j about the future realization of x. These idiosyncratic information sets evolve *via* SL.

Models with heterogeneous expectations are subject to the curse of dimensionality. The latter typically emerges for dynamic programming problems in large dimension, such as those with agent-level equations. In particular, a J-agent model implies a linear projection over state variables based on matrix of size higher than $J \times J$. Therefore the computational burden grows exponentially in J. In order to provide a tractable solution, without any sacrifice of the microfoundations of the model, we rely on the following main ingredients:

- **Nominal rigidities follow a Rotemberg pricing scheme.** Under this assumption, the first-order condition of the firms' maximization problem results in a recursive expression without using the law of iterated projections that is necessary in presence of menu costs à la Calvo.

- **Nonseparable utility function.** Another challenge in heterogeneous-agent models is the mapping between demand and labor supplies in the presence of potentially heterogeneous labor supplies. A non-separable utility function offers the advantage of muting the wealth effect on the labor supply, which ensures that all members of the family supply the same amount of hours and circumvents the aforementioned issue.

- **A government tax on price dispersion.** As common in the NK literature, we use a government tax on price dispersion to correct for the distortion generated by the heterogeneity in firms' prices and, hence, profits due to their heterogeneous inflation expectations.

- **Constant return to scale.** This hypothesis ensures that the hourly wages are homogeneous across agents as soon as they have the same productivity, even in case of heterogeneous labor supplies resulting from heterogeneous expectations.[11]

3.1.2 Households

Each agent j decides about their consumption, labor and saving plans in order to maximize the household's welfare:

$$\mathbb{E}_{j,t}\sum_{\tau=0}^{\infty}\beta^{\tau}u(c_{j,t+\tau},h_{j,t+\tau})=\frac{1}{J}\sum_{j=1}^{J}\mathbb{E}_{jt}\sum_{\tau=0}^{\infty}\beta^{\tau}\frac{1}{1-\sigma'}\left(c_{j,t+\tau}-\chi\frac{h_{j,t+\tau}^{1+\varphi}}{1+\varphi}\right)^{1-\sigma'}. \tag{20}$$

Utility $u(\cdot)$ is increasing in consumption $c_{j,t}$ and decreasing in labor $h_{j,t}$ where σ' and φ are curvature parameters and β is the discount factor. The non-separable utility function is based on the GHH utility function of Greenwood et al. [1988].

Agents face an intertemporal problem: they determine the value of their consumption $c_{j,t}$, hours worked $h_{j,t}$ and real bond holdings $b_{j,t}$ so as to maximize the welfare of the family under the following budget constraint which binds in every period:

$$c_{j,t}+b_{j,t}=\frac{i_{t-1}}{\pi_t}\frac{b_{j,t-1}}{\exp\left(\varsigma_g\hat{g}_t\right)}+w_th_{j,t}+T_{j,t}+z_{j,t}+\Pi_{j,t}, \tag{21}$$

where w_t is the real wage (assumed to be symmetric across all agents because they all have the same marginal product of labor under constant returns to scale); i_{t-1} the nominal interest rate payable on nominal bond holdings; π_t the inflation rate between periods $t-1$ and t; $\Pi_{j,t}$ the share of agent j of the real profits from monopolistic competition (see

[11] This assumption is not restrictive. We could address the issue of wage heterogeneity by introducing a labor packer that homogenizes labor demand across firms for instance.

Section 3.1.3); $T_{j,t}$ lump-sum government transfers that may be positive or negative; $z_{j,t}$ zero-sum intra-household transfers[12] and $\hat{g}_t$ an exogenous source of aggregate fluctuations (referred to as the risk-premium shock Smets and Wouters 2007) and affected by the elasticity parameter $\varsigma = \sigma' \frac{(1-\chi)}{\vartheta}$ which normalizes the shock with respect to the formulation in the textbook three-equation NK model.

Agents choose their consumption and savings plans conditional on their inflation and output-gap expectations. Hence, heterogeneity in expectations may entail heterogeneous consumption and wealth values, which poses a challenge for aggregation, in particular when saving is used in the production through capital goods. In the textbook NK model, labor is the only input in the production function and under the usual hypothesis of no restriction to the access to asset markets, the level of consumption is solely determined by the Euler equation and the idiosyncratic bond holdings do not affect current consumption. Therefore, the usual permanent income hypothesis holds at the agent-level. Note that the idiosyncratic saving stock measured by $b_{j,t}$ may be positive if savings or negative if borrowing, and that the family member is not allowed to run a Ponzi scheme, namely:

$$\lim_{T\to\infty} \mathbb{E}_{j,t}\left(\Lambda_{j,t,T}\frac{B_T}{P_T}\right) \geq 0$$

with $\Lambda_{j,t,T} \equiv \beta\frac{\partial u(c_{j,T},h_{j,T})/\partial c_{j,T}}{\partial u(c_{j,t},h_{j,t})/\partial c_{j,t}}$ the stochastic discount factor of the household. The aggregate demand for government bonds reads as follows: $B_t = \sum_{j=1}^{J} b_{jt}$.

Each household j solves the following problem:

$$\max_{\{c_{j,t},h_{j,t},b_{jt,}\}} \frac{1}{J}\sum_{j=1}^{J}\mathbb{E}_{jt}\sum_{\tau=0}^{\infty}\beta^{\tau}\left[\frac{1}{1-\sigma'}\left(c_{jt+\tau}-\chi\frac{h_{j,t+\tau}^{1+\varphi}}{1+\varphi}\right)^{1-\sigma'}\right.$$
$$\left.+\lambda'_{j,t+\tau}\left(\frac{i_{t-1+\tau}}{\pi_{t+\tau}}\frac{b_{j,t-1+\tau}}{\exp(\varsigma_g\hat{g}_{t+\tau})}+w_{t+\tau}h_{j,t+\tau}+z_{j,t+\tau}+\Pi_{j,t+\tau}-c_{j,t+\tau}-b_{j,t+\tau}\right)\right].$$

[12] These transfers impose homogeneous end-of-period wealth across households through an intra-risk sharing plan (see, e.g., Andrade et al. 2019). In each period, after making possibly heterogeneous consumption choices, each member's bond holding is made identical *via* an agreement to distribute the amount of bonds equally across the family. In detail, a transfer plan $z_{j,t} = b_{j,t} - B_t/J$ to each member j in each period t equalizes post-transfer wealth, even though consumption levels may differ across agents. In equilibrium, the sum of transfers is zero $\sum_{j=1}^{J} z_{j,t} = 0$.

The first-order conditions are given by:

$$w_t \lambda'_{j,t} = \chi h_{j,t}^{\varphi} \left(c_{j,t} - \chi \frac{h_{j,t}^{1+\varphi}}{1+\varphi} \right)^{-\sigma'},$$

$$\lambda'_{j,t} = \left(c_{j,t} - \chi \frac{h_{j,t}^{1+\varphi}}{1+\varphi} \right)^{-\sigma'},$$

$$\exp\left(\varsigma_g \hat{g}_t\right) \lambda'_{j,t} = i_t \beta \mathbb{E}_{jt} \frac{\lambda^c_{j,t+1}}{\pi_{t+1}}.$$

Log-linearizing each first-order condition yields:

$$\hat{w}_t = \varphi \hat{h}_{j,t}, \tag{22}$$

and:

$$\hat{\lambda}'_{j,t} = -\sigma' \left(\bar{c}_j - \chi \frac{\bar{h}_j^{1+\varphi}}{1+\varphi} \right)^{-1} \left(\bar{c}_j \hat{c}_{j,t} - \chi \bar{h}_j^{1+\varphi} \hat{h}_{j,t} \right), \tag{23}$$

$$\hat{\lambda}'_{j,t} = \hat{\imath}_t - \varsigma_g \hat{g}_t + \mathbb{E}_{jt} \left\{ \hat{\lambda}'_{j,t+1} - \pi_{t+1} \right\}. \tag{24}$$

where variables with a hat denote deviations from their deterministic steady-state values denoted by a bar. Eq. (22) shows that wages equal the marginal product of labor which is the same across all agents j due to their non-separable preferences. Moreover, at the deterministic steady state of the economy, all agents have the same information[13] and, therefore, consume the same amount of goods. It follows that: $\bar{c}_j = \bar{c}$, $\bar{h}_j = \bar{h}$, $\forall j$.

Equalizing Eqs. (23) and (24) yields:

$$\bar{c}\hat{c}_{j,t} - \chi \bar{h}^{1+\varphi} \hat{h}_{jt} = -\frac{\vartheta}{\sigma'} \left(\hat{\imath}_t - \mathbb{E}_{jt} \hat{\pi}_{t+1} \right) + \frac{\vartheta \varsigma_g}{\sigma'} \hat{g}_t + \mathbb{E}_{jt} \left[\left(\bar{c}\hat{c}_{jt+1} - \chi \bar{h}^{1+\varphi} \hat{h}_{jt+1} \right) \right], \tag{25}$$

with $\vartheta = \bar{c}_j - \chi \frac{\bar{h}^{1+\varphi}}{1+\varphi}$.

3.1.3 Firms

To introduce a monopolistic-competition framework, the production process of goods is divided between two types of firms: intermediate and final firms. Intermediate firms produce different types of goods which are imperfect substitutes. We assume that each member j owns an intermediate-sector firm j that produces an intermediate good y_j and generates the profit $\Pi_{j,t}$ (in Eq. (21)). Final firms produce a homogeneous good by combining all intermediate goods $\{y_j\}$, $j = 1, \ldots, J$.

Final sector. The final-good producers are retailers. They buy the intermediate

[13] At the deterministic steady-state, all agents hold steady-state beliefs and $\bar{a_j}^{\pi} = \bar{a_j}^{y} = 0$.

goods and package them into the aggregate supply of goods which, in equilibrium, equals the aggregate good demand from households, denoted by Y_t^D. On a perfectly competitive market, final producers take the price P of the goods as given and maximize profits as follows:

$$P_t Y_t^D - \sum_{j=1}^{J} p_{j,t} y_{j,t}, \tag{26}$$

subject to the supply constraint:

$$Y_t^D = \left(J^{-1/\epsilon} \sum_{j=1}^{J} {y_{j,t}}^{(\epsilon-1)/\epsilon} \right)^{\epsilon/(\epsilon-1)}, \tag{27}$$

which is the counterpart of the well-known Dixit-Stiglitz index with a finite amount of firms (the same holds for the price index (32) below). This supply constraint implies that the final-good producers have a technology which aggregates non-perfectly substitutable goods. This imperfect substitutability between all types of varieties j introduces monopolistic competition on the intermediate good market. Each good j is an imperfect substitute of degree $\epsilon > 1$, allowing intermediate firms to gain positive profits through a gap between their selling and producing prices. The intensity of the monopolistic competition is driven by $\epsilon/(\epsilon - 1)$, which is the mark-up over the marginal costs of intermediate firms.

The optimization problem of the final-good producers reads as follows:

$$L = P_t Y_t^D - \sum_{j=1}^{J} p_{j,t} y_{j,t} + \varrho_t \left[J^{-1/\epsilon} \sum_{j=1}^{J} {y_{j,t}}^{(\epsilon-1)/\epsilon} - \left(Y_t^D\right)^{(\epsilon-1)/\epsilon} \right]. \tag{28}$$

The associated first-order conditions are given by:

$$P_t - \varrho_t(\epsilon - 1)/\epsilon \left(Y_t^D\right)^{-1/\epsilon} = 0, \tag{29}$$

$$-p_{j,t} + \varrho_t(\epsilon - 1)/\epsilon J^{-1/\epsilon} {y_{j,t}}^{(-1)/\epsilon} = 0, \tag{30}$$

which can be rewritten as the standard CES downward-sloping demand function per firm j:

$$y_{j,t} = \frac{Y_t^D}{J} \left(\frac{p_{j,t}}{P_t} \right)^{-\epsilon}. \tag{31}$$

The aggregate-price index is given by:

$$P_t = \left[\frac{1}{J} \sum_{j=1}^{J} p_{j,t}^{1-\epsilon} \right]^{1/(1-\epsilon)}. \tag{32}$$

Log-linearizing Eqs. (31) and (32) leads to the following expressions:

$$\hat{y}_{j,t} = \hat{Y}_t^D - \epsilon \left(\hat{p}_{j,t} - \hat{P}_t \right), \tag{33}$$

$$\hat{P}_t = \frac{1}{J} \sum_{j=1}^{J} \hat{p}_{j,t}. \tag{34}$$

Expressing Eq. (34) in growth rates provides the expression for the inflation rate:

$$\hat{\pi}_t = \frac{1}{J} \sum_{j=1}^{J} \hat{\pi}_{j,t} \tag{35}$$

Intermediate sector. Each firm j in the intermediate sector has a linear production technology:

$$y_{j,t} = h_{j,t}, \tag{36}$$

where $y_{j,t}$ is their production and $h_{j,t}$ is their labor input.

Intermediate-goods producers solve a two-stage problem. In the first stage, taken the labor price w_t as given, firms hire labor $h^d_{j,t}$ in a perfectly competitive labor market in order to minimize their costs subject to the production constraint (36).

Stage 1: The first stage can be expressed as a profit-maximization problem:

$$\max_{\{y_{j,t} h^d_{j,t}\}} mc_{j,t} y_{j,t} - w_t h^d_{j,t} + \lambda_t \left[h^d_{j,t} - y_{j,t} \right], \tag{37}$$

where $mc_{j,t}$ denotes the real marginal cost of producing one additional good. The first-order condition leads to the expression of the real marginal cost:

$$mc_{j,t} = mc_t = w_t. \tag{38}$$

Because households exhibit the same labor productivity, all firms hire households at the same wage rate w_t. Once the firms have determined their marginal cost, the next step is to determine their mark-up over this marginal cost mc_t from the imperfect substitution of the good varieties.

Stage 2: In the second-stage problem, the firms operate under a Rotemberg price-setting mechanism. We define the Rotemberg price adjustment cost by:

$$ac_{jt} = \frac{\xi'}{2} \left(\frac{p_{jt}}{p_{jt-1}} - \bar{\pi} \right)^2 \frac{Y_t^D}{J}, \tag{39}$$

where $\xi' > 0$ is the price stickiness parameter, $\frac{Y_t^D}{J}$ is the average market share, $\bar{\pi}$ is the CB target.
The profit maximization becomes dynamic because of the adjustment costs over prices. In a monopolistic-competition setting, firms face the following individual demand for goods: $y_{j,t} = (p_{j,t}/P_t)^{-\epsilon} Y_t^D/J$. The problem faced by firms is then given by:

$$\max_{\{p_{j,t}\}} \mathbb{E}_{jt} \sum_{\tau=0}^{\infty} \Lambda_{j,t,t+\tau} \left((1 - \iota_{j,t+\tau})\, y_{j,t+\tau} \frac{p_{j,t+\tau}}{P_{t+\tau}} - e^{\varsigma_u \hat{u}_{t+\tau}} mc_{t+\tau} y_{j,t+\tau} - ac_{j,t+\tau} \right), \tag{40}$$

where $\Lambda_{j,t,t+\tau} = \beta^\tau \lambda'_{t+\tau}/\lambda'_t$ corresponds to the discount factor of agent j as previously defined in Eq. (3.1.2); $p_{j,t}$ is the individual price set by firm j, P_t is the aggregate price which sets the problem of firms in real terms. Variable $\hat{u}_t$ is an exogenous cost-push shock that captures exogenous changes in the cost structure of the firms. Parameter ς_u normalizes to one the parameter on this shock in the log-linearized form of the aggregate-supply equation (see Eq. (67) below). Note that the tax rate on the added value, $\iota_{j,t}$, is typically used in the NK literature to offset some market distortions and simplify the analysis of optimal policy. In the present case, given the presence of heterogeneity with respect to the benchmark textbook model, we assume that this tax is set by the government so as to offset the effect of the relative price dispersion on profits, i.e. $\iota_{j,t} = \frac{p_{j,t} - P_t}{p_{j,t}}$.

Plugging in the demand function $y_{j,t} = (p_{j,t}/P_t)^{-\epsilon} Y_t/J$, the objective function of the firms becomes:

$$\max_{\{p_{j,t}\}} \mathbb{E}_{jt} \sum_{\tau=0}^{\infty} \Lambda_{j,t,t+\tau} \left((1 - \iota_{j,t+\tau}) \left(\frac{p_{j,t+\tau}}{P_{t+\tau}} \right)^{1-\epsilon} \frac{Y_{t+\tau}^D}{J} - e^{\varsigma_u \hat{u}_{t+\tau}} mc_{t+\tau} \left(\frac{p_{j,t+\tau}}{P_{t+\tau}} \right)^{-\epsilon} \frac{Y_{t+\tau}^D}{J} - ac_{j,t+\tau} \right). \tag{41}$$

The first-order condition reads as:

$$(1 - \iota_{j,t}) \frac{(1-\epsilon)}{p_{j,t}} \left(\frac{p_{j,t}}{P_t} \right)^{1-\epsilon} \frac{Y_t^D}{J} + e^{\varsigma_u \hat{u}_t} \epsilon \frac{mc_t}{p_{j,t}} \left(\frac{p_{j,t}}{P_t} \right)^{-\epsilon} \frac{Y_t^D}{J} - \frac{\xi'}{p_{j,t-1}} \left(\frac{p_{j,t}}{p_{j,t-1}} - \bar{\pi} \right) \frac{Y_t^D}{J}$$
$$+ \mathbb{E}_{jt} \Lambda_{j,t,t+1} \frac{p_{j,t+1}}{p_{j,t}^2} \xi' \left(\frac{p_{j,t+1}}{p_{j,t}} - \bar{\pi} \right) \frac{Y_{t+1}^D}{J} = 0. \tag{42}$$

Log-linearizing this expression results in:

$$(\epsilon - 1)\, \hat{p}_{j,t} + (\epsilon - 1)/\epsilon \left(\varsigma_u \hat{u}_t + \widehat{mc}_t - \epsilon \hat{p}_{j,t} \right) = \frac{\xi'}{\epsilon} \bar{\pi} \hat{\pi}_{j,t} - \frac{\xi'}{\epsilon} \bar{\pi} \beta \mathbb{E}_{jt} \hat{\pi}_{j,t+1}, \tag{43}$$

with $\bar{mc} = (\epsilon - 1)/\epsilon$.

Rearranging terms leads to the final inflation equation for each producer j:

$$\hat{\pi}_{jt} = (\epsilon - 1) \frac{\varphi}{\bar{\pi} \xi'} \hat{h}_t + \beta \mathbb{E}_{jt} \hat{\pi}_{jt+1} + \hat{u}_t. \tag{44}$$

Note that setting $\varsigma_u \equiv \xi' \bar{\pi} / (\epsilon - 1)$ normalizes the shock in the linear equation because the marginal cost is the same across firms, $\widehat{mc}_t = \hat{w}_t$, and across households, $\hat{w}_t = \varphi \hat{h}_{jt}$ (from the wage-setting equation (22)).

3.1.4 Authorities

Monetary policy. The monetary policy authority, namely the CB, sets the nominal interest rate i as a function of the deviation of inflation from its target and of the output gap. We use here a general formulation of the monetary policy that allows for interest-rate smoothing and history-dependent inflation targets. We can then easily contrast the effects of various configurations of monetary policy under SL and RE in Section 5.

Precisely, the monetary policy rule reads as:

$$i_t = i_{t-1}^{\rho} \times \left(\bar{\imath} \times \left(\frac{\tau_t}{\bar{\tau}} \right)^{\phi_\pi} \times \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y} \right)^{1-\rho}, \tag{45}$$

where parameter $\rho \in [0, 1]$ reflects a concern for interest-rate smoothing, $\{\phi_\pi, \phi_y\}$ are the two reaction coefficients to, respectively, inflation and output gaps and τ_t is the CB target that reads as:

$$\tau_t = \pi_t \tau_{t-1}^{(1-\omega)}, \tag{46}$$

where $0 \leq \omega \leq 1$ is a discounting parameter; see Budianto et al. [2023]. This formulation of the inflation target nests the two polar cases of inflation targeting (IT, hereafter), which corresponds to $\omega = 1 \Leftrightarrow \tau_t = \pi_t$ and price-level targeting (PLT, hereafter), where $\omega = 0$ and $\tau_t = \pi_t \tau_{t-1} = \pi_t P_{t-1} = P_t$. For any intermediary values of ω, the CB implements average-inflation targeting (AIT, hereafter) with a geometric discounting of past inflation gaps.

The log-linearization of the monetary policy rule and the inflation target yields:

$$\hat{\imath}_t = \rho \hat{\imath}_{t-1} + (1 - \rho) (\phi_\pi \hat{\tau}_t + \phi_y \hat{y}_t), \tag{47}$$

and

$$\hat{\tau}_t = \hat{\pi}_t + (1 - \omega) \hat{\tau}_{t-1}. \tag{48}$$

Government. The government implements the tax $\iota_{j,t}$ on the added-value of firms and borrow B_t from the households, while the expenditure side includes interest payments and lump-sum transfers (where it is assumed that all household members receive the same amount). The budget constraint of the government is then:

$$\sum_{j=1}^{J} \iota_{j,t} y_{jt} + B_t = \sum_{j=1}^{J} T_{jt} + B_{t-1} i_{t-1} / \pi_t. \tag{49}$$

3.1.5 Equilibrium conditions

Intermediate sector. The equilibrium in the intermediate-good market is given by:

$$\sum_{j=1}^{J} y_{j,t} = \frac{Y_t^D}{J} \sum_{j=1}^{J} \left(\frac{p_{j,t}}{P_t} \right)^{-\epsilon}, \tag{50}$$

which can be log-linearized as:

$$J\bar{y} \sum_{j=1}^{J} \hat{y}_{j,t} = \bar{Y}^D \left[J\hat{y}_t^D - \epsilon \left(\sum_{j=1}^{J} \hat{p}_{j,t} - J\hat{P}_t \right) \right]. \tag{51}$$

Using the definition in Eq. (34) rules out the effect of price dispersion on output and allows one to rewrite the equilibrium on the goods market as:

$$\frac{1}{J} \sum_{j=1}^{J} \hat{y}_{j,t} = \hat{y}_t^D, \tag{52}$$

where $J\ \bar{y} = \bar{Y}^D$.

Final goods sector. The resource constraint is given by:

$$Y_t^D = \sum_{j=1}^{J} \left(c_{j,t} + \frac{\xi'}{2} \left(\pi_{j,t} - \bar{\pi} \right)^2 \frac{Y_t^D}{J} \right), \tag{53}$$

where the second term on the right-hand side correspond to the adjustment costs stemming from the Rotemberg's price adjustment mechanism.

Log-linearizing Eq. (53) yields:

$$\hat{y}_t^D = \frac{1}{J} \sum_{j=1}^{J} \hat{c}_{j,t}, \tag{54}$$

with $\bar{c} = \bar{y} = \bar{Y}^D / J$.

Labor market. Equilibrium in the labor market is reached when the aggregate labor demand from firms satisfies:

$$\sum_{j=1}^{J} h_{j,t}^d = \sum_{j=1}^{J} h_{j,t}, \tag{55}$$

or, after log-linearization:

$$\sum_{j=1}^{J} \hat{h}_{j,t}^d = \sum_{j=1}^{J} \hat{h}_{j,t}, \tag{56}$$

because firms and households share the same steady-state values for the hours worked.

3.1.6 Aggregation

Labor demands dispersion. In presence of non-separable preferences, the labor supply in Eq. (22) is the same across households:

$$\hat{h}_{j,t} = \hat{h}_t. \tag{57}$$

Combining Eq. (36) with the firm-specific demand for intermediate inputs (33) leads to the following expression:

$$\hat{h}^d_{j,t} = \hat{y}_{j,t} = \hat{Y}^D_t - \epsilon \left(\hat{p}_{j,t} - \hat{P}_t\right). \tag{58}$$

One can note that if a firm has a pricing strategy that is different from the average, i.e. $\hat{p}_{j,t} \neq \hat{P}_t$, their labor demand differs from the aggregate labor demand, i.e. $\hat{y}_{j,t} \neq \hat{Y}^D_t$, which implies a dispersion in both labor demands and output across firms. This dispersion in labor demand needs to be matched with the homogeneous labor supplies from the households in Eq. (57). To do so, we use the following assumption:

> **Assumption 1.** *To map potentially heterogeneous labor demands with homogeneous labor supplies, households evenly split their working hours across all firms at no cost, which translates into:*
>
> $$h_{j,t} = \sum_{j=1}^{J} \frac{h^d_{j,t}}{J} \tag{59}$$

Log-linearizing Eq. (59) and plugging in Eq. (58) allows one to express any agent-specific change in labor supply to change in aggregate demand:

$$\begin{aligned}
\hat{h}_{j,t} &= \sum_{j=1}^{J} \frac{\hat{h}^d_{j,t}}{J} \\
&= \sum_{j=1}^{J} \frac{\hat{Y}^D_t - \epsilon \left(\hat{p}_{j,t} - \hat{P}_t\right)}{J} \\
&= \hat{Y}^D_t + \epsilon \hat{P}_t - \epsilon \sum_{j=1}^{J} \frac{(\hat{p}_{j,t})}{J} \\
&= \hat{Y}^D_t
\end{aligned} \tag{60}$$

Hence, even in presence of heterogeneous labor demands induced by price dispersion, the equilibrium in the labor market in Eq. (56) holds.

Aggregate demand. Consider the agent-level log-linearized Euler equation (25) and aggregate across all household members as:

$$\sum_{j=1}^{J}\left[\bar{c}_j\hat{c}_{j,t}-\chi\bar{h}^{1+\varphi}\hat{h}_{j,t}\right]=\sum_{j=1}^{J}\left[-\frac{\vartheta}{\sigma'}\left[\hat{\imath}_t-\mathbb{E}_{jt}\hat{\pi}_{j,t+1}\right]+\mathbb{E}_{j,t}\left(\bar{c}_j\hat{c}_{j,t+1}-\chi\bar{h}^{1+\varphi}\hat{h}_{j,t+1}\right)+\frac{\vartheta\varsigma}{\sigma'}\hat{g}_t\right],$$

which may be rearranged into:

$$\bar{c}_j\hat{c}_t-\chi\bar{h}^{1+\varphi}\hat{h}_t=-\frac{\vartheta}{\sigma'}\hat{\imath}_t+\mathbb{E}_t\left(\bar{c}_j\hat{c}_{t+1}-\chi\bar{h}^{1+\varphi}\hat{h}_{t+1}+\frac{\vartheta}{\sigma'}\hat{\pi}_{t+1}\right)+\frac{\vartheta\varsigma}{\sigma'}\hat{g}_t,$$

where $\mathbb{E}_t\hat{c}_{t+1}=\frac{1}{J}\sum_{j=1}^{J}\mathbb{E}_{jt}\hat{c}_{j,t+1}$, $\mathbb{E}_t\hat{\pi}_{t+1}=\frac{1}{J}\sum_{j=1}^{J}\mathbb{E}_{jt}\hat{\pi}_{j,t+1}$ and $\mathbb{E}_t\hat{h}_{t+1}=\frac{1}{J}\sum_{j=1}^{J}\mathbb{E}_{j,t}\hat{h}_{j,t+1}$.

Equilibrium in labor market allows one to write $\hat{h}_t=\hat{h}_t^d=\hat{y}_t$, while equilibrium in the intermediate-good market entails $\hat{y}_t=\hat{y}_t^D=\hat{c}_t$. The previous condition reads as:

$$\left(\bar{c}-\chi\bar{h}^{1+\varphi}\right)\hat{y}_t=-\frac{\vartheta}{\sigma'}\hat{\imath}_t+\mathbb{E}_t\left(\left(\bar{c}-\chi\bar{h}^{1+\varphi}\right)\hat{y}_{t+1}+\frac{\vartheta}{\sigma'}\hat{\pi}_{t+1}\right)+\frac{\vartheta\varsigma}{\sigma'}\hat{g}_t,$$

Recall that at the steady state, the hours worked are normalized to one, thus $\bar{c}-\chi\bar{h}^{1+\varphi}=1-\chi$ and $\vartheta=1-\chi/1+\varphi$. Recall also that $\varsigma=\sigma'\frac{(1-\chi)}{\vartheta}$. Hence, the aggregate Euler equation in a compact form reads as:

$$(1-\chi)\,\hat{y}_t=-\frac{\vartheta}{\sigma'}\hat{\imath}_t+\mathbb{E}_t\left((1-\chi)\,\hat{y}_{t+1}+\frac{\vartheta}{\sigma'}\hat{\pi}_{t+1}\right)+(1-\chi)\,\hat{g}_t. \tag{61}$$

Aggregate supply. From the micro-level NK Phillips curve (44), we may aggregate over agents j:

$$\sum_{j=1}^{J}\hat{\pi}_{jt}=\sum_{j=1}^{J}\left[(\epsilon-1)\frac{\varphi}{\bar{\pi}\xi'}\hat{h}_t+\beta\mathbb{E}_{jt}^{*}\hat{\pi}_{jt+1}+\hat{u}_t\right],$$

which becomes:

$$\hat{\pi}_t=(\epsilon-1)\frac{\varphi}{\bar{\pi}\xi'}\hat{h}_t+\beta\mathbb{E}_t\hat{\pi}_{t+1}+\hat{u}_t. \tag{62}$$

Note that we can replace aggregate labor by aggregate output as in the usual textbook formulation, which results in:

$$\hat{\pi}_t=(\epsilon-1)\frac{\varphi}{\bar{\pi}\xi'}\hat{y}_t+\beta\mathbb{E}_t\hat{\pi}_{t+1}+\hat{u}_t. \tag{63}$$

The monetary policy rule (47) is left unchanged.

3.1.7 Convergence between separable and non-separable utilities

In this section, we show which restrictions on the parameters of the non-separable utility function allow the model to correspond to the one derived from separable preferences and reconcile the form of our three-equation NK model with heterogeneous expectations with the usual textbook formulation of NK models under RE (see, e.g., Galí 2015, Chap. 3).

Marginal utility of consumption. Let λ_t and λ_t' denote respectively the marginal utility of consumption under separable and non-separable preferences respectively. These are given by:

$$\hat{\lambda}_t = -\sigma \hat{c}_t \quad \text{and} \quad \hat{\lambda}_t' = -\sigma' \left(\frac{1-\chi}{1-\chi/\left(1+\varphi\right)} \right) \hat{c}_t$$

Imposing $\hat{\lambda}_t = \hat{\lambda}_t'$ results in the condition on σ' under which both models exhibit the same marginal utilities of consumption:

$$\sigma' = \frac{1-\chi/\left(1+\varphi\right)}{1-\chi}\sigma. \tag{64}$$

We may substitute σ' by σ into (61) as follows:

$$\hat{y}_t = \mathbb{E}_t \left\{ \hat{y}_{t+1} \right\} - \frac{1}{\sigma}\mathbb{E}_t \left\{ \hat{\imath}_t - \hat{\pi}_{t+1} \right\} + \hat{g}_t. \tag{65}$$

Slope of the New Keynesian Phillips Curve. Nonseparability in utility also affects the real-wage setting, and thus the marginal cost and the slope of the NK Phillips curve:

$$\kappa = \left(\epsilon - 1\right) \frac{\left(\sigma + \varphi\right)}{\bar{\pi}\xi} \text{ and } \kappa' = \left(\epsilon - 1\right) \frac{\varphi}{\bar{\pi}\xi'}.$$

By imposing $\kappa = \kappa'$, we derive ξ':

$$\xi' = \frac{\varphi}{\left(\sigma + \varphi\right)}\xi. \tag{66}$$

Under this second condition, the aggregate supply curve may be rewritten as follows:

$$\hat{\pi}_t = \left(\epsilon - 1\right) \frac{\left(\sigma + \varphi\right)}{\bar{\pi}\xi}\hat{y}_t + \beta \mathbb{E}_t \hat{\pi}_{t+1} + \hat{u}_t \tag{67}$$

4 Simulating social learning in Dynare

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].

4.1 Simulations under SL

The toolbox is seamlessly implemented using the usual Dynare's syntax. Users may write their SL model as they otherwise would with a RE model in a `.mod` file. They need to specify the parameter values relative to the model and the SL algorithm as well as the shock sequence. They may also choose whether the SL news shocks (or mutations) are drawn from a predefined seed or not. The only remaining requirement to solve and simulate a model under SL is to first call sequentially the two functions `SL_get_pq` and then `SL_chain` after spelling out the model. This is necessary because the function `SL_get_pq` solves for matrices P, Q and R (see Eq. (18)) and the function `SL_chain` then simulates the model under RE and SL, given these matrices and the chain of exogenous shocks provided. Our toolbox allows one to simulate impulse response functions (IRFs) and stochastic time series of the model as is commonly used in the standard RE macroeconomic literature.

The toolbox files provided alongside this paper[14] contain extensive comments. We provide there the code to implement the model of Section 3.1 and derive the results from Section 5 below. Appendix A also displays illustrations of a run of the model.

4.2 Calibration

We use standard parameter values because the purpose of the exercise hereafter is to illustrate the use of the toolbox rather than 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. We can use Eq. (64) to calculate the non-separable case which yields $\sigma' = 2.5$. The elasticity of substitution is set to $\varepsilon = 4$ (or, equivalently, the markup is set at 33%) and the slope of the Phillips curve is set to $\kappa = 0.05$. From Eq. (66) and by setting the inverse of the Frish elasticity to $\varphi = 1$, we can see that these parameter values correspond to adjustment costs of $\xi' = 59.7015$ in the non-separable case. The benchmark monetary policy parameters represent a standard inflation targeting rule à la Taylor [1993] without interest-rate smoothing $\{\phi_\pi, \phi_y, \rho, \omega\} = \{1.5, 0.125, 0, 1\}$. This parameter set fulfils the Taylor principle and ensures determinacy under RE in the baseline model.[15] The yearly inflation target (or trend) is set to 2%. We use $\{\rho_g, \rho_u\} = \{0.8, 0.5\}$, $\sigma_\varepsilon^g = 0.1$ and $\sigma_\varepsilon^g = 0.05$.

As for the SL parameter values, we follow, *inter alia*, Arifovic et al. [2018, 2023] and set the mutation probabilities (which correspond to the frequencies of individual news reception) to $\{\mu_\pi, \mu_y\} = \{0.3, 0.3\}$ and the memory in the discounted past history process in Eqs. (6) to $\{\rho_\pi, \rho_y\} = \{0.8, 0.8\}$. The noise of the news shocks in Eq. (5) is set to $\sigma_\zeta = 0.3$ for both inflation and output gap and we use $J = 300$ agents.

For the purpose of illustration, a representative run of the model under SL and RE is displayed in Appendix A and simulated moments are reported in Table 1.

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

[15] In Subsection 5.2 where we explore the monetary policy parameter space, we only consider values that ensure a determinate model under RE.

5 Results

We illustrate the added value of the micro-founded SL framework, our solution method and the associated `Dynare` toolbox with an application to optimal monetary policy design under SL and RE. We first derive a welfare criterion, then discuss the trade-off between inflation and output gap variance under a standard IT regime and conclude with optimal history-dependent monetary policy rules.

5.1 Aggregate welfare approximation

The micro-foundations of the SL model allow us to derive the second-order approximation of the welfare criterion necessary to assess optimal monetary policy. We decompose welfare under SL into two components. The first component is associated with the volatility of the aggregate variables and corresponds to welfare in the RE framework. The second component arises under SL as a result of the heterogeneity in expectations and is associated with the volatility of the idiosyncratic variables. The detailed derivations and explicit forms are deferred to Appendix B.

We express the utility of an average SL agent j as

$$E_j\left(U_{j,t}\right) \simeq u_0 - u_{\gamma\gamma} E_j var_t\left(\hat{\gamma}_{j,t}\right) - u_{yy} var_t\left(\hat{y}_t\right) - u_{\rho\rho} E_j var_t\left(\hat{\rho}_{j,t}\right) - u_{\pi\pi} E_j var_t\left(\hat{\pi}_{j,t}\right), \quad (68)$$

where 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 corresponding welfare function is the discounted sum of the utility of the average SL agents given by

$$\mathcal{W}_t = \frac{E_t E_j\left(U_{j,t}\right)}{1-\beta}. \quad (69)$$

It is straightforward to notice that under SL, both macroeconomic volatility and heterogeneity among agents reduce the aggregate welfare of households. By contrast, in the absence of heterogeneity across agents in the RE model, $var\left(\hat{\gamma}_{j,t}\right) = var_t\left(\hat{\rho}_{j,t}\right) = 0$ and $E_j var_t\left(\hat{\pi}_{j,t}\right) = var\left(\hat{\pi}_t\right)$ and the utility function boils down to a decreasing function of aggregate output and inflation volatility. Hence, SL implies a welfare loss with respect to the RE counterpart due to the dispersion among individual consumption and price levels stemming from heterogeneous inflation and output gap expectations.

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. (68).

5.2 Stabilization trade-off under IT

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 $\rho = 0$ and $\omega = 1$, which reduces the monetary policy rule (45) 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.

5.2.1 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 gaps. In what follows, we analyze the variance of inflation and output gaps 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 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. 5 provides another look at this trade-off. It displays the so-called IT efficiency frontier under RE (solid line) and SL (dashed line), that is the constellation 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:[16]

$$\mathcal{L}(\alpha) = \alpha Var(\hat{\pi}_t) + (1 - \alpha) Var(\hat{y}_t). \tag{70}$$

The closer the frontier is to the origin, the more efficient monetary policy is 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).

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

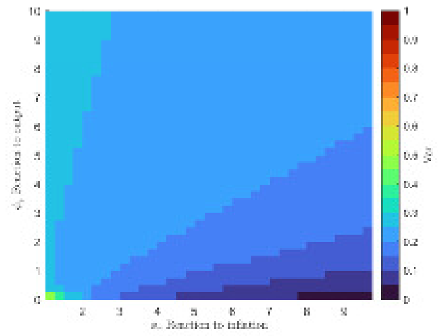

(a) Mean inflation variance RE $var\left(\hat{\pi}_t\right)$

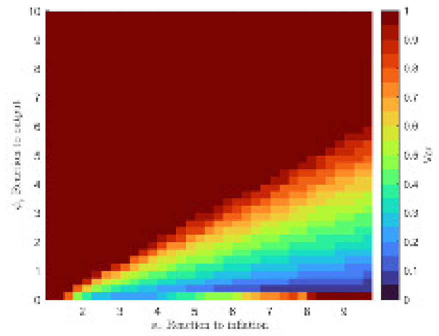

(b) Mean inflation variance SL $var\left(\hat{\pi}_t\right)$

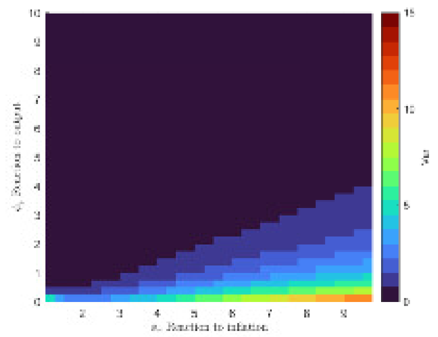

(c) Output variance RE $var\left(\hat{y}_t\right)$

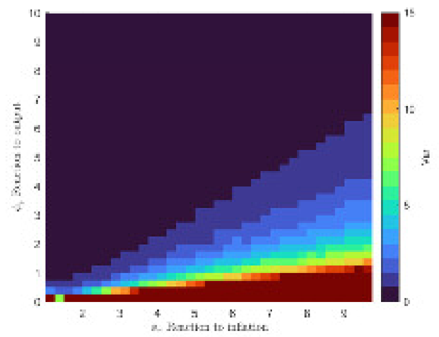

(d) Output variance SL var $\left(\hat{y}_t\right)$

<u>Notes:</u> We use a discrete grid of 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 1000 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

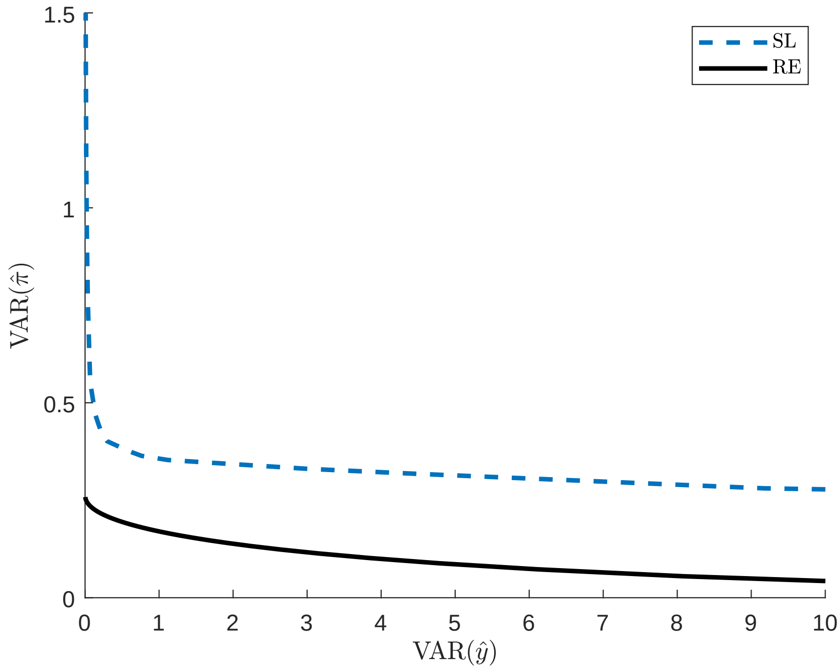

Figure 2: Efficiency Frontiers

The main takeaway from Fig. 5 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 asymptot under SL. This divergence of outcomes under the two expectational 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 affects inflation, more than output gap, because inflation is almost self-fulfilling in the NK model.[17]

An additional major difference between RE and SL arises when the CB prioritizes stabilizing output (on the left-hand side of Fig. 5). In this region, the SL frontier is much steeper than the RE frontier and the two diverge strongly. 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 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].

[17] This is particularly the case because β is close to one and the NKPC is fairly flat.

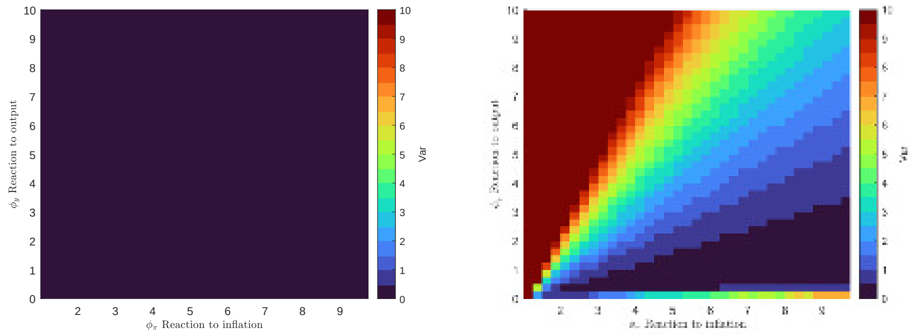

(a) Dispersion of the idiosyncratic relative prices ($\hat{\rho}_{j,t}$) under RE

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

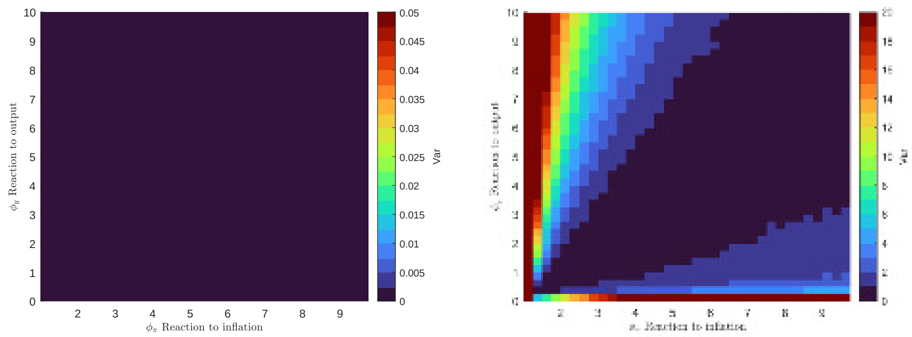

(c) Dispersion of the idiosyncratic consumption levels ($\hat{\gamma}_{j,t}$) under RE

(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

5.2.2 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 (68), 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 SL (right-hand side panel) as a function of the two IT parameters ϕ_π (x-axis) and ϕ_y (y-axis).

Of course, the RE values of these idiosyncratic welfare components are always zero. 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 is, the more disagreement between agents and the more dispersed the prices they set. Intuitively, more volatile inflation implies that the innovation process of the SL algorithm (i.e. the mutations) may generate non-steady state expectations that temporarily survive the tournament process and turn self-fulfilling because they happen to generate the lowest forecast errors in the population of forecasts. More volatile inflation implies larger and more frequent shifts in the average past inflation trend that the agents learn to forecast through SL. 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 micro-foundations of the model with SL expectations. Using Eqs. (25), (57) and (58), we can rewrite the consumption of agent j as:

$$\widehat{c}_{j,t} = \mathbb{E}_{j,t}\widehat{y}_{t+1} - \sigma^{-1}(\widehat{\imath}_t - \mathbb{E}_{j,t}\widehat{\pi}_{t+1}) + \widehat{g}_t. \tag{71}$$

Similarly, using Eq. (44), their idiosyncratic inflation rate is equivalent to:

$$\widehat{\pi}_{j,t} = \beta\mathbb{E}_{j,t}\widehat{\pi}_{t+1} + \kappa\widehat{y}_t + \widehat{u}_t. \tag{72}$$

From Eq. (72), we immediately can 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. (71) 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 to determine 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 increase 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 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 results 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.

5.2.3 Welfare implications

Aggregating now these different components into the welfare criterion, Fig. 4 shows the welfare outcomes in the parameter space $\{\phi_\pi, \phi_y, \rho, \omega\} = \{(1, 10], (0, 10], 0, 1\}$ 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 welfare differences between the two expectation schemes are mostly in levels: in both cases, the high welfare zone corresponds to a strong reaction to inflation and a more moderate reaction to output gap. Due to the combination of both the welfare cost of aggregate output volatility and the welfare cost of the dispersion of idiosyncratic consumption and price levels, some output stabilization is beneficial under SL – to see that look at the particularly low welfare levels obtained at the bottom right corner of Fig. 4b where inflation reaction in the Taylor rule is strong but output reaction is absent. 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).

5.3 Optimal policy results

Empirical monetary policy rules may also include a lag to account for interest rate smoothing over time. Additionally, over the recent years, history-dependent policy rules have gained interest. For instance, in August 2020, the Fed announced that the inflation target will be 'an average over time.' To account for these refinements, in what follows we consider optimal Taylor rules over the full set of parameter values of Eqs. (45)-

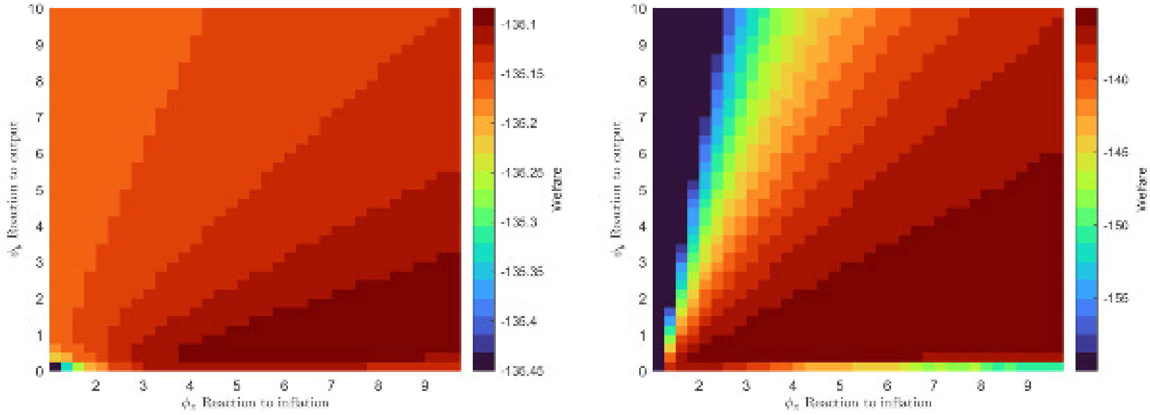

(a) Welfare under RE (b) Welfare under SL

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

Figure 4: Welfare under IT

(48). Specifically, we use $\{\phi_\pi, \phi_y, \rho, \omega\} = \{(1, 10], (0, 10], [0, 1], (0, 1]\}$. In the toolbox, the function `OSR_sl` finds these parameter values that maximize $\mathcal{W}_t$.[18]

Under RE, the optimization problem is fairly simple and implies finding the combination of aggregate output and inflation variances that minimizes $\mathcal{W}_t$. Under SL, the problem is more complex because it also requires to take into account the adverse welfare effects of the dispersion in the individual prices and consumption levels. An additional layer of complexity stems from the idiosyncratic random draws and the non-linear feedback loops between the SL expectation formation process and the realization of the endogenous variables. For this reason, it is advisable to use a global optimization algorithm.[19]

Table 1 contrasts the simulated moments of the model under the standard IT policy previously considered in Section 5.2, the optimal policy under RE, and the optimal policy under SL. The parameter allowing for history-dependence in the target ω and the interest-rate smoothing parameter ρ introduce interesting time dependence in the interest rate setting. It is worth recalling that the textbook RE model under IT (which corresponds to $\omega = 1$ and $\rho = 0$) does not contain any lagged state variables (such as those stemming from consumption habits or price indexation in more sophisticated versions of the NK model). In this baseline version, only when persistence is introduced in the monetary policy rule – either *via* interest rate smoothing (i.e. $\rho > 0$) or *via* history-dependent inflation target (i.e. $\omega < 1$) – are inflation and output gaps no longer 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

[18] The function can also be set to minimize $\Theta(\alpha)$

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

	(I)	(II)	(III)
***Monetary policy regime*:**	***Benchmark IT***	***Optimal RE pol.***	***Optimal SL pol.***
Calibration $\{\phi_\pi, \phi_y, \rho, \omega\}$	$\{1.5, 0.125, 0, 1\}$	$\{10, 3.36, 0, 0.01\}$	$\{3.91, 2.76, 0.04, 0.06\}$
Inflation variance $var\,(\hat{\pi}_t)$:			
RE model	0.25801	0.065189	0.099459
SL model	0.90274	0.11281	0.1578
Output gap variance $var\,(\hat{\pi}_{j,t})$:			
RE model	1.637	2.3719	1.0626
SL model	2.8969	13.144	5.3164
Variance of individual consumption levels $var\,(\hat{\gamma}_{j,t})$:			
RE model	0	0	0
SL model	1.3215	0.5115	0.21003
Variance of individual price dispersion $var\,(\hat{\rho}_{j,t})$:			
RE model	0	0	0
SL model	0.9694	0.1565	0.17584
Welfare $(\mathcal{W}_t)$:			
RE model	-135.2	-135.02	-135.04
SL model	-136.24	-135.33	-135.23
Loss in permanent consumption equivalent w.r.t the deterministic steady state			
RE model	0.085%	0.031%	0.035%
SL model	0.41%	0.13%	0.10%

Notes: See Fig. 1.

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

forecast and the tournament selection. Therefore, under SL, inflation and output gaps instead exhibit persistence no matter the specification of monetary policy.

In Table 1, Col. I, we see that all aggregate moments under a standard IT regime exhibit a larger variance with SL expectations than with RE. Welfare under SL further decreases with respect to the RE counterpart as a result of losses resulting from dispersion in individual prices and consumption levels. Consequently, welfare under SL is considerably smaller than under RE.

Looking into optimal policy design, we find the common result that PLT is optimal under RE (Col. II). The corresponding set of policy parameters that maximize welfare is $\{\phi_\pi, \phi_y, \rho, \omega\} = \{10, 3.36, 0, 0.01\}$, which refers to the most aggressive PLT rule[20] without any interest-rate smoothing and some reaction to the output gap. 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. This lower inflation variance further implies lower price dispersion among SL firms via less heterogeneous inflation expectations.

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

[20] Recall that we restrict our parameter space to $\phi_\pi = 10$, which is large for empirical standards.

in the efficiency frontiers 5 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. (71)). Interestingly, this second effect dominates, and consumption patterns are less heterogeneous under the optimal RE policy than under the calibrated IT rule. Taken together, these observations result in the optimal policy under RE being welfare-improving under both expectation schemes, not only under RE.

Finally, the optimal policy under SL is reported in Col. III of Table 1 and corresponds to a fairly aggressive average-inflation targeting rule with a long history (a small ω-value), a fairly strong reaction to output gap and a particularly mild interest-rate smoothing, namely $\{\phi_\pi, \phi_y, \rho, \omega\} = \{3.91, 2.76, 0.04, 0.06\}$. Under this rule, 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 a far better output stabilization and less dispersed consumption patterns under the SL optimal rule (the variance of the output gap is more than twice smaller under the SL than under the RE optimal 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 note that under backward-looking expectations, and adaptive learning in particular, history-dependent rules have been found to generate *more* volatility than standard IT frameworks (see Honkapohja and McClung 2023, Amano et al. 2020). A key result of this paper is instead that optimal policies are robust to different expectation specifications: optimal policies under RE and under SL are welfare-improving in both models, no matter whether expectations are rational or result from social interactions. This result is good news for policymakers that 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.

6 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 and a welfare criterion that incorporates both the costs of aggregate volatility and dispersion in idiosyncratic pricing and consumption decisions. We further deliver a 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 a fair amount of history-dependence in the inflation target is 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 genetic algorithms 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 particularly promising mechanism to simulate so-called animal spirits and belief contagions in asset markets in larger-scale macroeconomic models.

References

Klaus Adam and Albert Marcet. Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory*, 146(3):1224–1252, 2011. 9

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. 33

Philippe Andrade, Gaetano Gaballo, Eric Mengus, and Benoit Mojon. Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics*, 11(3):1–29, 2019. 2, 13

George-Marios Angeletos, Fabrice Collard, and Harris Dellas. Quantifying confidence. *Econometrica*, 86(5):1689–1726, 2018. 4

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

Jasmina Arifovic, James Bullard, and Olena Kostyshyna. Social learning and monetary policy rules. *The Economic Journal*, 123(2001):38–76, 2013. 3, 9, 11

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. 3, 9, 23

Jasmina Arifovic, Alex Grimaud, Isabelle Salle, and Gauthier Vermandel. Social Learning and Monetary Policy at the Effective Lower Bound. Technical report, Bank of Canada; revised 2023, 2023. 2, 3, 23

Florin O Bilbiie. The new keynesian cross. *Journal of Monetary Economics*, 114:90–108, 2020. 2

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. 2

Flora Budianto, Taisuke Nakata, and Sebastian Schmidt. Average inflation targeting and the interest rate lower bound. *European Economic Review*, 152:104384, 2023. 18

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. 2

Jordi Galí. *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press, 2015. 22

Jeremy Greenwood, Zvi Hercowitz, and Gregory W Huffman. Investment, capacity utilization, and the real business cycle. *American Economic Review*, pages 402–417, 1988. 12

J. H. Holland. *Adaptation in Natural and Artificial Systems*. Ann arbor: University of michigan press edition, 1975. 2

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

Seppo Honkapohja and Nigel McClung. On robustness of average inflation targeting. *Available at SSRN 4021712*, 2023. 33

The MathWorks Inc. Matlab version: 9.13.0 (r2022b), 2022. URL `https://www.mathworks.com`. 22, 31

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. 3, 22

Greg Kaplan, Benjamin Moll, and Giovanni L Violante. Monetary policy according to hank. *American Economic Review*, 108(3):697–743, 2018. 2

Guido Lorenzoni. Optimal monetary policy with uncertain fundamentals and dispersed information. *Review of Economic Studies*, 77(1):305–338, 2010. 10

Domenico Massaro. Heterogeneous expectations in monetary DSGE models. *Journal of Economic Dynamics and Control*, 37(3):680–692, 2013. ISSN 01651889. doi: 10.1016/j.jedc.2012.11.001. 2

Athanasios Orphanides and John C. Williams. Inflation scares and forecast-based monetary policy. *Review of Economic Dynamics*, 8(2):498–527, 2005. 27

Pablo Ottonello and Thomas Winberry. Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, 88(6):2473–2502, 2020. 2

Michael Reiter. Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665, 2009. 2

Frank Smets and Rafael Wouters. Shocks and frictions In US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97:586–606, 2007. 13

John B Taylor. Discretion versus policy rules in practice. 39:195–214, 1993. 23

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. 2

A Representative run of the model

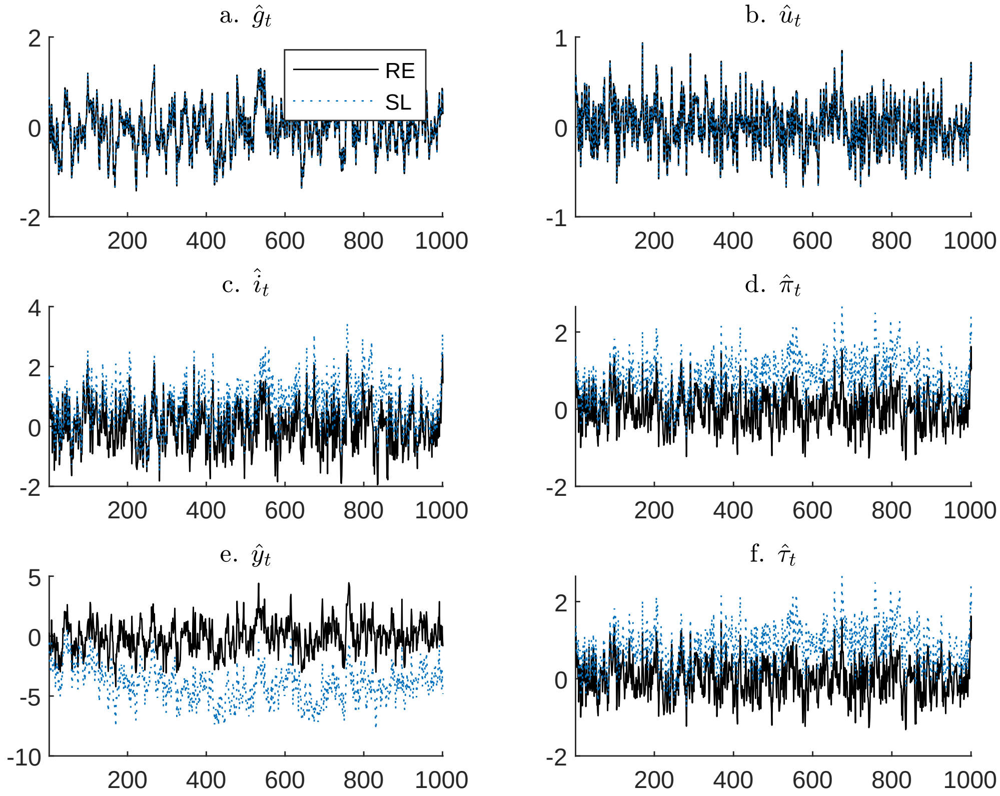

Notes: We removed the first 300 periods of initialization of the SL model.

Figure 5: Representative run of the model under SL and RE (note the de-anchoring and re-anchoring of the SL model during the simulation)

B Welfare function

The welfare function is the discounted sum of the average utility across family members:

$$\mathcal{W}_t = \frac{1}{J}\sum_{j=1}^{J} U_{jt} + \beta \mathcal{W}_{t+1},$$

where the individual utility function of each household member reads as:

$$U_{jt} = \frac{1}{1-\sigma'}\left(c_{jt} - \chi \frac{h_{jt}^{1+\varphi}}{1+\varphi}\right)^{1-\sigma\prime}.$$

Few changes are needed to be able to express the agent-specific utility function in the same aggregate-variable terms as seen in macro textbooks. First, recall that all household members supply the same amount of labor, $h_{jt} = Y_t/J$. This means we can rewrite h_{jt} in term of aggregate output. Second, recall that family members enjoy different consumption levels, which can be expressed in terms of fractions of aggregate consumption as $\gamma_{jt} = Nc_{jt}/C_t$, with $C_t = \sum_{j=1}^{J} c_{jt}$. The resulting utility of agent j is given by:

$$U_{jt} = \frac{1}{1-\sigma'}\left(\gamma_{jt}\frac{C_t}{J} - \chi\frac{(Y_t/J)^{1+\varphi}}{1+\varphi}\right)^{1-\sigma\prime}.$$

To take into account the role of nominal rigidities, we use the resource constraint in Equation 53 to replace aggregate output as follows: $C_t = Y_t^D\left(1 - \frac{1}{J}\sum_{j=1}^{J}\frac{\xi'}{2}\left(\pi_{jt} - \bar{\pi}\right)^2\right)$. In addition, we express output in supply terms using the general-equilibrium condition in the intermediate sector in Equation 50: $Y_t^D = Y_t\left(\frac{1}{J}\sum_{j=1}^{J}\left(\frac{p_{jt}}{P_t}\right)^{-\epsilon}\right)^{-1}$. This last term corresponds to the dispersion across firms' prices that entails an output loss.

The utility of the j-th household member, expressed in terms of aggregate output and inflation, is given by:

$$U_{jt} = \frac{1}{1-\sigma'}\left(\gamma_{jt}y_t\frac{\left(1 - \frac{1}{J}\sum_{j=1}^{J}\frac{\xi'}{2}\left(\pi_{jt} - \bar{\pi}\right)^2\right)}{\left(\frac{1}{J}\sum_{j=1}^{J}\rho_{jt}^{-\epsilon}\right)} - \chi\frac{y_t^{1+\varphi}}{1+\varphi}\right)^{1-\sigma\prime},$$

where $\rho_{jt} = p_{jt}/P_t$ and $y_t = Y_t/J$.

Using a Taylor expansion of the utility function up to second-order terms, abstracting from co-variance terms, the welfare function can be expressed in terms of asymptotic moments, namely unconditional mean $E_t(\cdot)$ and variance $var_t(\cdot)$. It leads to the following expression:

$$E_t\left(U_{jt}\right) \simeq \bar{U}_j + \frac{U_{\gamma\gamma}}{2}var_t\left(\gamma_{jt}\right) + \frac{U_{yy}}{2}var_t\left(y_t\right) + J\frac{U_{\rho\rho}}{2}E_jvar_t\left(\rho_{jt}\right) + J\frac{U_{\pi\pi}}{2}E_jvar_t\left(\pi_{jt}\right),$$

where derivatives are computed from a symbolic toolbox.

Note that $E_jvar_t\left(\pi_{jt}\right)$ is the average volatility of inflation across firms, while $E_jvar_t\left(\rho_{jt}\right)$ denotes the average variation in relative prices across firms. Note that in a representative-agent model, the first term is the average inflation rate while the second term is zero, as the relative price across producers is always one. The mean utility can be expressed in logs to be consistent with the model's definition, as follows:

$$E_t\left(U_{jt}\right) \simeq \bar{U}_j + \frac{U_{\gamma\gamma}}{2\bar{\gamma}^2}var_t\left(\hat{\gamma}_{jt}\right) + \frac{U_{yy}}{2\bar{Y}^2}var_t\left(\hat{y}_t\right) + J\frac{U_{\rho\rho}}{2}E_jvar_t\left(\hat{\rho}_{jt}\right) + J\frac{U_{\pi\pi}}{2\bar{\pi}^2}E_jvar_t\left(\hat{\pi}_{jt}\right),$$

where $\bar{\gamma} = \bar{y} = 1$.

In sum, the average welfare for the planner reads as:

$$E_t E_j (U_{jt}) \simeq \frac{1}{J} \sum_{j=1}^{J} \bar{U}_{jt}, \tag{73}$$

which can be expressed as:

$$E_t E_j (U_{jt}) \simeq \bar{U}_j + \frac{U_{\gamma\gamma}}{2} E_j var_t (\hat{\gamma}_{jt}) + \frac{U_{yy}}{2} var_t (\hat{y}_t) \tag{74}$$

$$+ J \frac{U_{\rho\rho}}{2} E_j var_t (\hat{\rho}_{jt}) + J \frac{U_{\pi\pi}}{2\bar{\pi}^2} E_j var_t (\hat{\pi}_{jt}). \tag{75}$$

Average utility is connected to average welfare as follows:

$$\mathcal{W}_t = \frac{E_t E_j (U_{jt})}{1-\beta}, \tag{76}$$

Rational expectation & representative-agent utility. Note that, in absence of heterogeneity across agents (such as under RE), there is no change in $\hat{\gamma}_{jt}$ and $\hat{\rho}_{jt}$ across time, thus $var(\hat{\gamma}_{jt}) = var(\hat{\rho}_{jt}) = 0$. In addition, there is no difference across inflation rates $E_j var(\hat{\pi}_{jt}) = var(\hat{\pi}_t)$. The utility function therefore becomes:

$$E_t (U_t) \simeq \bar{U} + \frac{U_{yy}}{2} var_t (\hat{y}_t) + J \frac{U_{\pi\pi}}{2\bar{\pi}^2} var_t (\hat{\pi}_t). \tag{77}$$