# Social Learning and Monetary Policy at the Effective Lower Bound*

Jasmina ARIFOVIC[1] Alex GRIMAUD[2] Isabelle SALLE[3]
Gauthier VERMANDEL[4]

2023-03-23

**Abstract**

This paper develops a model that jointly accounts for the missing disinflation in the wake of the Great Recession and the subsequently observed inflation-less recovery. The key mechanism works through heterogeneous expectations that may durably lose their anchoring to the central bank (CB)'s target and coordinate on particularly persistent below-target paths. The welfare cost associated with persistent low inflation may be reduced if the CB announces to the agents its target or its own inflation forecasts, as communication helps coordinate expectations. However, the CB may lose its credibility whenever its announcements become decoupled from actual inflation.

*Keywords*: Inflation targeting, Heterogeneous expectations, Effective lower bound, Communication.

*JEL codes*: E32; E52; E70; C82.

*The present work has benefited from fruitful discussions and helpful suggestions from Jason Allen, Philippe Andrade, Jim Bullard, Fabrice Collard, Ben Craig, Pablo Cuba-Borda, George Evans, Jordi Galí, Tomáš Holub, Cars Hommes, Seppo Honkapohja, Eleni Iliopulos, Robert Kollmann, Douglas Laxton, Albert Marcet, Domenico Massaro, Bruce McGough, Athanasios Orphanides, Fabien Tripier, Jonathan Witmer and Raf Wouters. We also thank the participants of the Workshop on Adaptive Learning on May 7-8, 2018, in Bilbao, the 2018 EEA-ESEM meeting in Cologne, the seminar at the Bank of Latvia on March 23, 2019, the workshop on Expectations in Dynamic Macroeconomic Models at the 2019 GSE Summer Forum in Barcelona, the 2019 Computing Economics and Finance Conference in Ottawa, the 15th Dynare Conference in Lausanne, the 22nd Central Bank Macroeconomic Modeling Workshop in Dilijan, in particular our discussant Junior Maih, the internal seminar at the Bank of Canada on September 22, 2019, and the 2019 SNB research conference in Zurich and the online CEBRA conference on September 2-4, 2020. This work has received funding from the European Union Horizon 2020 program (Marie Sklodowska-Curie grant No.721846) and from the OeNB Anniversary fund (project No.18611). We thank William Ho for helpful research assistance. None of the above are responsible for potential errors in the paper.

[1]Department of Economics, Simon Fraser University, Burnaby, BC, CA. In January 2022, before the current revision of the paper, Jasmina tragically passed away. Therefore, while acknowledging the breadth of her contributions and her intellectual leadership, the views expressed in the current version of the paper can only be attributed to the remaining authors.

[2]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)

[3]Department of Economics, University of Ottawa (isalle@uottawa.ca), Amsterdam School of Economics, University of Amsterdam and Tinbergen Institute, NL

[4]École Polytechnique UMR 7641 & Université Paris-Dauphine UMR 8007, Paris FR (gauthier@vermandel.fr)

# 1 Introduction

This paper develops a micro-founded general-equilibrium model in which the dynamics of heterogeneous expectations provide a crucial and empirically-grounded ingredient to account for the business-cycle dynamics.

During this period in the US and Europe, historically low interest rates have accompanied persistently low but stable inflation and below-target inflation expectations, as depicted by Figure 1. Substantial changes in price levels have failed to materialize both in the wake of the downturn and thereafter, in what resembles an inflation-less recovery. This low-inflation narrative is hard to unfold within the standard New Keynesian (NK) models because those models generate implausible macroeconomic dynamics at the effective lower bound (ELB hereafter).

Under rational expectations (RE), the dynamics are indeterminate at the ELB (Benhabib et al. 2001), while under recursive learning – its most common alternative – deflationary spirals arise as soon as the ELB binds for long enough (Evans et al. 2008, Ozden & Wouters 2021). Yet, real-world expectations have remained in this indeterminate and unstable region of the inflation-output state space after the financial crisis, as depicted in Figure 2, without giving rise to any excessively volatile or diverging inflation. Thus, the data appears to defy the predictions of standard modelling frameworks.

Recent developments offer alternative models of stable dynamics at the ELB but remain within the framework of the representative agent; see, in particular, Gabaix (2020) for a parsimonious and micro-founded example. However, leaving out the heterogeneity of real-world expectations may be problematic: not only does this conception of expectations conflict with the extensive empirical evidence of unanchored and dispersed forecasts,[1] but it also

[1] See, *inter alia*, Mankiw et al. (2003) in survey data from professional forecasters; and Branch (2004)

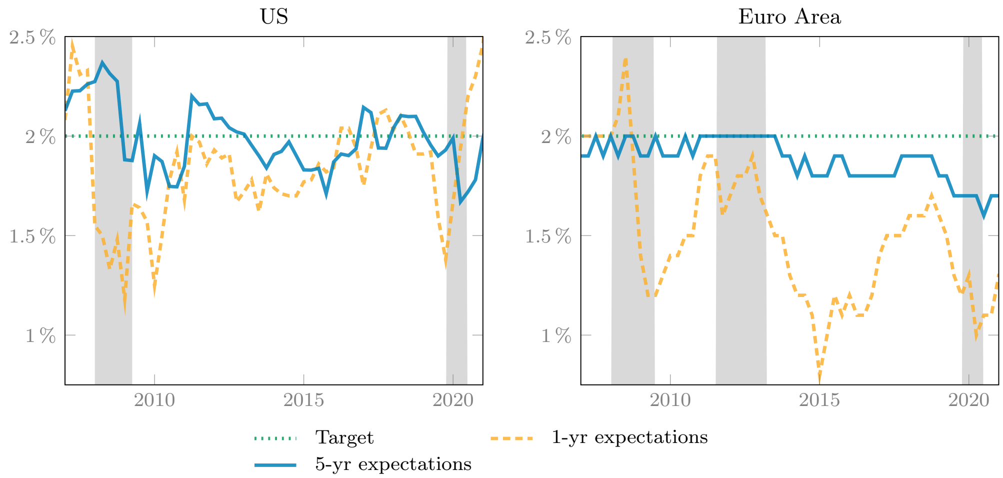


Notes: The shaded areas represent the recessions as dated by the National Bureau of Economic Research (NBER) and Centre for Economic Policy Research (CEPR), and the green dashed lines the inflation targets. Data is from the Survey of Professional Forecasters (SPF). For the US, we use the mean PCE from the Philadelphia Fed survey; for the Euro Area, the mean HICP from the European Central Bank (ECB) survey.

Figure 1: Inflation expectations in the US and Euro Area 2008–2021

does not leave any room for the coordinating role of CB communication. Heterogeneity in expectations is an essential challenge that CBs face when managing expectations. This is especially true in periods of high uncertainty and macroeconomic volatility, such as the aftermath of the COVID-19 pandemic, where disagreement between agents is likely to be reinforced.

There exist models of heterogeneous expectations, as detailed in the literature review below, but their avoidance of implausible dynamics at the ELB comes at the cost of ad-hoc and empirically questionable constraints and/or deep sophistication.[2] For instance, Ozden

from households. Coibion et al. (2019) show that more than half of the surveyed firms and households do not know the value of the Fed inflation target. One-year-ahead household inflation expectations are on average 1.5 percentage points (p.p.) above the target, and the cross-sectional dispersion reaches up to 3 p.p. (Coibion et al. 2020). Moreover, Cornea-Madeira et al. (2019) estimate a model with evolutionary selection among forecasting heuristics on survey and inflation data and find considerable heterogeneity in forecasting models that is also time-varying. For laboratory evidence in this respect, see, e.g., the survey by Hommes (2021).

[2]In learning models in general, imposing an ad-hoc floor on deflation also rules out these explosive paths (Evans et al. 2022).

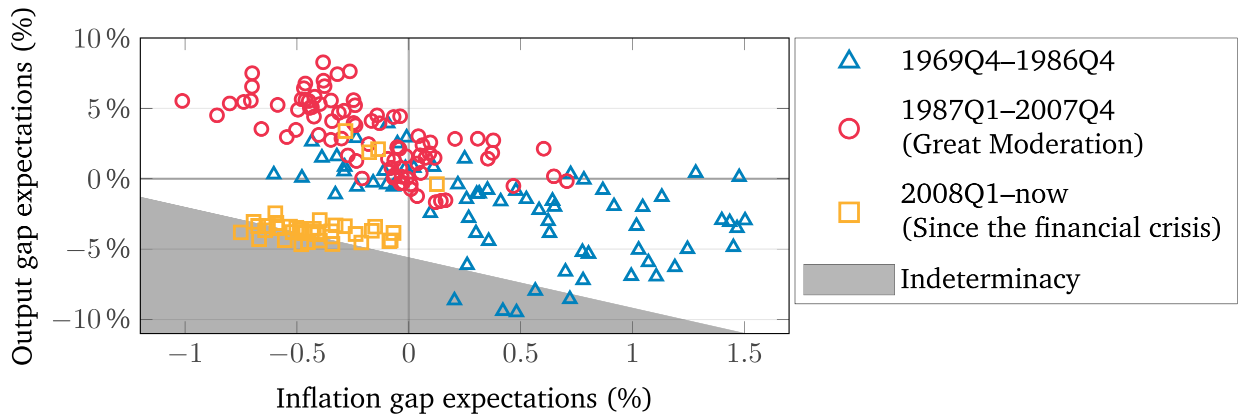


Notes: The shaded area denotes the region of the state space that violates the Blanchard-Kahn determinacy condition under RE and leads to diverging deflationary spirals under recursive learning. The white area denotes the determinate region of the state space, which is also the basin of attraction of the target under recursive learning. Data on expectations are taken from the SPF. The output gap is computed using a linear trend. Calibration of the NK model is taken from Galí (2015).

Figure 2: (Ir)relevance of the New Keynesian model with rational expectations since the ‘new normal’

(2021) shows that a sufficiently high proportion of the agents need to keep their expectations anchored at the target even during an ELB episode so as to limit the diverging trend and induce convergence back to the steady state. In Andrade et al. (2019), heterogeneity only exists at the ELB and pertains to the length of the ELB episode, while inflation and output forecasts remain homogeneous.

We therefore contribute to the literature by developing a model that combines the appealing properties from, e.g. Gabaix (2020) – in particular parsimony and micro-foundations and the emergence of stable dynamics at the ELB – with time-varying heterogeneity in expectations. The model can be matched against moments from real-world expectations and restores the coordinating role of CB communication. In our framework, the heterogeneous expectations are the engine that endogenously drives the economy into stable and extended periods at the ELB and allows our model to account for the recent economic experiences.

We develop the micro-foundations of a NK model with heterogeneous expectations that evolve using a parsimonious evolutionary learning process and nest the RE homogeneous-agent benchmark. Learning dynamics create room for expectations to be persistently off target and play an autonomous role in driving business cycles, so that recessive episodes - and ELB episodes - need not be generated by exogenous and persistent technology or financial shocks.[3] In our model, agents employ steady-state learning; i.e. they form beliefs about the long-run values of inflation and output, which easily translates into the issue of expectation anchoring.

Specifically, we have chosen to model these dynamics through the use of a social learning (SL) process. Our choice is motivated by the parsimony of this class of learning models, their ability to match experimental findings and the evolutionary role of heterogeneity in the adaptation of the agents. In these models, agents collectively adapt to an ever-changing environment in which their own expectations contribute to shape the macroeconomic variables that they are trying to forecast. This feature is well suited to self-referential economic systems such as standard macroeconomic models. SL expectations also find an intuitive interpretation that is reminiscent of the idea of epidemiological expectations, in which 'expert forecasts' only gradually diffuse across the entire population (Carroll 2003).

In a novel effort within the related literature,[4] we take our stylized model to the data and show that it is able to *jointly* replicate ten salient business cycle moments from the Survey of Professional Forecasters (SPF) and the main US macroeconomic time series. This moments include the frequency of ELB episodes, major dimensions of heterogeneity in expectations

[3] Angeletos et al. (2018) investigate the role of strategic uncertainty in the presence of heterogeneous information within a general-equilibrium model. However, those authors use a real business cycle (RBC) model, which implies that monetary policy is left out.

[4] Del Negro & Eusepi (2011) attempt to replicate expectation data with RE models. Milani (2007) fits an adaptive learning NK model to macroeconomic time series only. Closer to our contribution, Slobodyan & Wouters (2012a,b) estimate an NK model on both macroeconomic and expectation times series. However, the authors use exogenous autocorrelated shocks on expectations to reproduce the observed persistence in the data.

and a substantial share of the persistence in output and inflation data. This empirical exercise is already a remarkable result given the parsimony of the model. Our empirical application makes two contributions to the literature: (*i*) a moment-matching routine for a non-linear model under heterogeneous expectations and (*ii*) data-consistent values to the learning parameters for which there are no observable counterparts.

A second major contribution is to show that our model *endogenously* produces stable dynamics at the ELB. Those stable dynamics correspond to recent inflation-less recoveries. We can loosely define such a recovery as one in which inflation persists for an extended period of time below its target, *i.e.* the ELB binds, but output expands. The model matches particularly well the probability of the ELB on nominal interest rates to bind despite the relatively modest amplitude and i.i.d. structure of the fundamental shocks. Those ELB episodes are not the result of large exogenous shocks but are an endogenous product of the interplay between learning and the small i.i.d. shocks.

In our model, once agents have coordinated on pessimistic expectations, the transition back to the target can be particularly long; expectations have become unanchored and, per their self-fulfilling nature, sustain the bust. Hence, we offer a reading of the recent economic experience as resulting from a long-lasting coordination of agents on pessimistic expectations rather than persistent and exogenous shocks.

Given that our model nests the RE homogeneous-agent benchmark, we interpret the dispersion of expectations as a friction and quantify the ensuing welfare loss with respect to the RE outcome. We find that heterogeneous expectations entail a consumption loss of almost 0.87% with respect to the RE allocation. From there, a natural follow-up analysis is to introduce an additional monetary policy instrument, namely CB communication, and investigate whether it may offset the costs of forecast dispersion. To address this question, we exploit the flexibility of the SL model, which enables us to integrate CB communication into the learning process of the agents. From two simple communication examples, we show the the critical role CB credibility plays in the ability of its announcements to reshape expectations. The CB may lose credibility whenever the announcements become decoupled from the actual realizations of inflation. Moreover, accurate but below-target inflation forecasts may turn self-defeating by coordinating expectations on a pessimistic depression. In light of these observations, We then discuss the impact of these observations on recent policy debates, such as the forward-guidance puzzle or the adoption of (temporary) higher inflation targets.

**Related literature** Our paper primarily relates to the growing literature on DSGE models with heterogeneous expectations. The earlier contributions in this area ignore the issue of the ELB, focus instead on the equilibrium learnability and dynamic of the three-equation reduced-form NK model. For instance, De Grauwe (2011) models upward- and downward-biased beliefs in a three-equation NK model. In this paper, evolutionary selection based on a heuristic-switching model (HSM, hereafter)- as developed in the seminal contribution of Brock & Hommes (1997) – creates endogenous waves of pessimism and optimism that mimic business cycles. Once heterogeneous beliefs are accounted for, flexible inflation targeting helps stabilize inflation, a result confirmed in a laboratory experiment by Hommes et al. (2019).

Branch & McGough (2009) introduce a two-type expectation model in the NK reduced-form framework, where a fraction of the agents are rational and the remainder are adaptive. They find that heterogeneous expectations may lead to multiple equilibria. Gasteiger (2014) and Di Bartolomeo et al. (2016) characterize optimal policy in such a framework. Branch &

McGough (2010) extend this framework by allowing heterogeneity to be time-varying using an HSM. Anufriev et al. (2013) show that, when many types of biased expectations co-exist in a frictionless DSGE model, the Taylor principle is no longer sufficient to achieve the target.[5] Massaro (2013) focuses on the micro-foundations of the NK model with an HSM. Closer to the present paper, Arifovic et al. (2013) discuss the robustness of the Taylor principle to SL expectations, but ignore ELB dynamics and micro-foundations.

Investigating the interaction between the ELB and HSM, Busetti et al. (2017) develop a model in which agents switch between a misspecified forecasting rule that may unanchor their long-run expectations and a mean-reverting adaptive rule. Ozden (2021) extends their framework to consider regime-switching between an ELB and a normal environment with unconstrained monetary policy. This author shows that the ELB episodes need to be short in order to avoid divergence along a deflationary path that is prevalent under adaptive learning. This divergence problem, which our framework overcomes, is extensively discussed under learning and homogeneous agents by Ozden & Wouters (2021).

Our paper also adds to the literature on communication under learning. Earlier contributions emphasize that deviations from RE provide a strong rationale for CB communication (Orphanides & Williams 2007). The learning literature generally concludes that communication is stabilizing under learning in models where communication use a model-consistent forecasting model.[6] One important assumption in these models is that communication is fully credible.

By contrast, in our model, the CB's credibility evolves endogenously as a result of the realized inflation gap. The existing contributions that come the closest to our treatment of

[5]This result echoes the well-known earlier finding of Orphanides & Williams (2004) under adaptive learning, and stands in contrast with the conclusion of Gabaix (2020), in which boundedly rational agents, albeit homogeneous, imply a wider determinacy region than under RE.

[6]see, *inter alia*, Orphanides & Williams (2005, 2007), Eusepi & Preston (2010).

endogenous credibility come from Hachem & Wu (2017), Hommes & Lustenhouwer (2019a) and Goy et al. (2020). Hommes & Lustenhouwer (2019a) consider the stability properties of a model with an EBL and heterogeneous expectations. In this model, agents may anchor their expectations at the target or switch to follow past inflation, should the target be missed and the CB lose credibility. In a similar vein, Goy et al. (2020) consider a model in which the agents switch between either integrating the CB forward-guidance announcements in their expectations or forming adaptive expectations. Hachem & Wu (2017) abstract from the NK setup but integrate social dynamics in the evolution of the CB's credibility. They make a case for credible gradual communication in order to disinflate the economy, and aggressive communication in order to reflate an economy trapped at the ELB.

Andrade et al. (2019) stand out thanks to their they focus on a different kind of heterogeneity. In their model, agents hold heterogeneous beliefs as to the length of the liquidity trap, which leads to both a pessimistic and an optimistic interpretation of the CB forward guidance as well as two distinct co-existing expectations regarding the length of the ELB episode. Nevertheless, heterogeneity in their setting is not dynamic and only prevails at the ELB.

Finally, there exist larger-scale and considerably more complex DSGE models than our present framework, in which non-linearities play a key role in accounting for the recent economic experience; see, *inter alia*, Gust et al. (2017), Lindé & Trabandt (2019). Our work is particularly related to the NK models with multiple equilibria where the persistent slump after the Great Recession is understood as an exogenously driven regime-switch from the targeted equilibrium to the deflationary steady state (Schmitt-Grohé & Uribe 2017, Aruoba et al. 2017, Arifovic et al. 2018, Lansing 2020). However, the coordination mechanism generating liquidity traps in our model is fundamentally different from the one used in the above-cited contributions.

In the context of our model, agents never coordinate on the low-inflation steady state, nor do they contemplate the possibility of a regime switching between the two steady states. The target is the only stable equilibrium under SL and expectations always remain within its basin of attraction, which is shown to be larger under SL than the determinacy region under RE. As a result of a series of adverse fundamental shocks, expectations may travel to regions of that basin from which convergence back to target takes a very long time; in these regions, the ELB binds and the pessimistic expectations are self-defeating per the self-fulfilling nature of the expectations.

To model expectations, we use a SL mechanism similar to Arifovic et al. (2013, 2018). Yet, our work differs substantially. Importantly, those two theoretical contributions study the asymptotic stability of the NK model under SL. Our focus is on the short-term fluctuations arising from the interplay between fundamental shocks and learning dynamics and the empirical performances of the SL. Arifovic et al. (2018) interpret liquidity trap episodes as the coordination of expectations on the low-inflation steady-state that is stable under their learning mechanism. Our agents have a finite memory and our empirical calibration differs from theirs, which does not allow us to generalize their result to our setup. In fact, in our model, the low-inflation state is unstable under SL as it belongs to the basin of attraction of the target; if expectations shift into the low-inflation state, they will eventually converge back to the target, but after a considerable amount of time.

The rest of the paper proceeds as follows. In Section 2, we develop the model; the moment-matching and calibration exercise is presented in Section 3; the dynamic properties of the model are analyzed in Section 4; Section 5 discusses CB communication; and Section 6 concludes.

# 2 The model

We first describe the building blocks of the model, then present the solution under the RE benchmark and finally focus on how expectations evolve under SL. The micro-foundations of the model under heterogeneous SL expectations are developed in Appendix A.1.

## 2.1 A piecewise linear New Keynesian model

Our model builds on the workhorse three-equation NK model into which we introduce heterogeneous expectations. The time and the number of agent types are discrete. There are $N$ agent-types, indexed by $j = 1, \ldots, N$, that only differ in terms of expectations but otherwise share the same characteristics (in particular in terms of preferences and technology). The main assumptions are welfare maximizing households with endogenous labor supply and nonseparable preferences, while firms are monopolistic competitors and face menu costs *à la* Rotemberg. In this setup, expectations result from idiosyncratic shocks (mutations) and social interactions (tournament); see Section 2.3. Appendix A.1 provides the explicit micro-foundations of the SL model and shows how aggregate heterogeneous expectations may be obtained from the arithmetic mean of the individual expectations among all $N$ agent types, indexed by $j = 1, ..., N$.

In what follows, we provide the underlying linearized equations of the micro-founded model. As in its textbook version, the NK model is summarized by three core equations. All variables below are expressed in terms of deviation from their steady-state levels, which in turn correspond to the CB target.

The first equation, the IS curve, describes aggregate demand:

$$\widehat{y}_t = \mathbb{E}_t^* \left\{ \widehat{y}_{t+1} \right\} - \sigma^{-1} (\widehat{\imath}_t - \mathbb{E}_t^* \left\{ \widehat{\pi}_{t+1} \right\}) + \widehat{g}_t, \tag{1}$$

where $\widehat{y}_t$ is the output gap, $\widehat{\imath}_t$ the nominal interest rate set by the CB, $\widehat{\pi}_t$ the deviation of the inflation rate from the target (hence, $\widehat{\imath}_t - \mathbb{E}_t^* \widehat{\pi}_{t+1}$ is the real expected interest rate), $\widehat{g}$ an exogenous real disturbance, $\sigma > 0$ the inter-temporal elasticity of substitution of consumption (based on a CRRA utility function), and $\mathbb{E}_t^*$ the (possibly boundedly rational) *aggregate* expectation operator based on information available at time $t$.

The second equation is the forward-looking NK Phillips curve that summarizes the supply side:

$$\widehat{\pi}_t = \beta \mathbb{E}_t^* \left\{ \widehat{\pi}_{t+1} \right\} + \kappa \widehat{y}_t + \widehat{u}_t, \tag{2}$$

where $0 < \beta < 1$ represents the discount factor, $\kappa > 0$ a composite parameter capturing the sensitivity of inflation to the output gap and $\widehat{u}_t$ an exogenous cost-push shock.

The third equation describes the law of motion of the nominal rate. Monetary policy implements a flexible inflation-targeting regime subject to the ELB constraint, which results in the following non-linear Taylor rule:

$$\widehat{\imath}_t = \max\{-\overline{r}; \phi^\pi \mathbb{E}_t^* \left\{ \widehat{\pi}_{t+1} \right\} + \phi^y \mathbb{E}_t^* \left\{ \widehat{y}_{t+1} \right\}\}, \tag{3}$$

where $\phi^\pi$ and $\phi^y$ are the reaction coefficients to the respective gaps in inflation and output, and $\overline{r} \equiv \pi^T + \rho$ the steady-state in level of interest rate associated with the inflation target $\pi^T$ and the households' discount rate $\rho \equiv -\log(\beta)$. The forward-looking rule highlights the emphasis of CBs on expectations as contemporaneous variables are not instantaneously observable.

We now present the solution of the model under the RE benchmark and then detail how the agents' expectations evolve under the SL process.

## 2.2 The model under rational expectations

In this section, we consider RE and impose $\mathbb{E}_t^*(\cdot) = E(\cdot \mid I_t)$ as the RE operator given the information set $I_t$ common to all agents in period $t$. We solve for the minimal state variable (MSV) solution using the method of undetermined coefficients. All details are provided in Appendix A.2.

It is well known that the ELB introduces a non-linearity in the Taylor rule and generates an additional deflationary steady-state (Benhabib et al. 2001). Hence, expressing the model in reduced form is complicated by this non-linearity, and we need to disentangle two pieces, one around the target and one where the ELB is binding.[7]

A short digression through the one-dimensional Fisherian model easily illustrates this configuration. Figure 3 displays inflation and interest rate dynamics, abstracting from the production side: the inflation target corresponds to $\widehat{\pi} = 0$ and the deflationary steady state to $\widehat{\pi}^{elb}$. Provided that $\widehat{\pi}^{elb} \leq 0 \leq \pi^T$, the two equilibria co-exist.

Coming back to the two-dimensional model, we have to specify a process for the exogenous shocks. In the rest of the paper, we consider white noise shocks only, so $\widehat{g}$ and $\widehat{u}$ are non-observable *i.i.d.* processes. In this case, the MSV solution boils down to a noisy constant without persistence. The presence of a floor on the nominal rate makes this solution piece-

[7]We follow here the related NK literature and impose the ELB constraint in the log-linearized model around the targeted steady state to describe the dynamics around the low inflation state, see, *inter alia*, Guerrieri & Iacoviello (2015). This method gives a second-best measure of the dynamics around the deflationary steady state. A first-best would be to log-linearize the model around this second steady state, but this would result in an MSV solution involving extra additional state variables (Ascari & Sbordone 2014) and, hence, additional coefficients to learn under SL (see Section 2.3). However, the benefits in terms of qualitative results are unlikely to outweigh the costs of such a complication of the learning process of the agents.

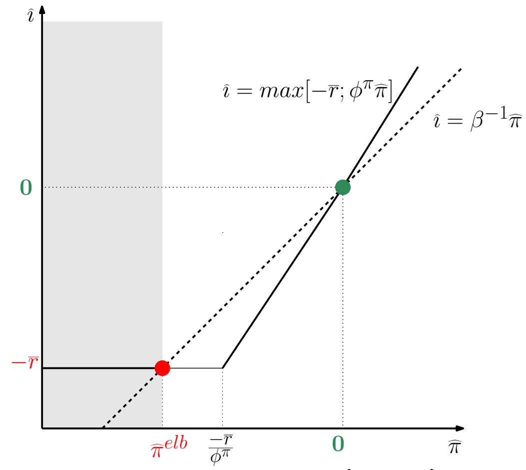


Notes: We can write the log-approximated Fisher equation as follows: $\widehat{\imath} = \beta^{-1}\widehat{\pi}$. At the targeted steady state (in green), no deviation occurs: $\widehat{\imath} = \beta^{-1}\widehat{\pi} = 0$. At the ELB (in red), we can derive an equilibrium such that $-\overline{r} = \beta^{-1}\widehat{\pi}^{elb} \Leftrightarrow \widehat{\pi}^{elb} = -\overline{r}\beta$. Provided that $\widehat{\pi}^{elb} \leq \pi^{T}$, the two equilibria co-exist. The shaded area is indeterminate under RE and unstable under adaptive learning (Evans et al. 2008).

Figure 3: Co-existence of two steady states under the ELB constraint

wise linear:

$$\widehat{z}_t = [\widehat{y}_t \ \widehat{\pi}_t]' = \begin{cases} a^T + \chi^g \widehat{g}_t + \chi^u \widehat{u}_t, \text{ if } i_t > 0 \\ a^{elb} + \chi^g \widehat{g}_t + \chi^u \widehat{u}_t, \text{ if } i_t = 0, \end{cases} \tag{4}$$

where the first case is the law of motion when the ELB is not binding (denoted by a ‘T’ superscript) and the second case when the ELB is binding (denoted by an ‘elb’ superscript). The exact expression of the matrix coefficients can be found in Appendix A.2. Note that as variables are expressed in terms of deviation from their steady-state values at the target, we have $a^T = (0 \ 0)'$. Additionally, under RE, the combination of white-noise shocks and a forward-looking rule implies that expectations remain anchored at the target and the ELB never binds. We now introduce expectations under SL.

## 2.3 Expectations under social learning

Under SL, we relax the assumption of homogeneous agents endowed with RE and consider instead a population of $N$ heterogeneous and interacting agents, indexed by $j = 1, \cdots, N$. We now define $\mathbb{E}_t^*(\cdot) = \mathbb{E}_{j,t}^{SL}(\cdot \mid I_{j,t})$ to be the expectation operator under SL given the information set $I_{j,t}$ available in period $t$ to agent $j$. The information set is agent-specific as it contains, in addition to the history of past inflation and output gaps up until period $t-1$, the individual's current and past forecasts, which need not be shared with the whole population.

**Individual forecasting rules** Following Arifovic et al. (2013, 2018), we assume that all agents are endowed with forecasting rules that are consistent with the MSV solution but involve agent-specific coefficients that they revise over time. As detailed in Appendix A.1.1, in any period $t$, each agent $j$ is therefore entirely described by a two-component forecast $[a_{j,t}^y, a_{j,t}^\pi]'$ and her expectations read as:

$$\mathbb{E}_{j,t}^{SL}\left\{\hat{z}_{j,t+1}\right\} = \begin{bmatrix} \mathbb{E}_{j,t}^{SL}\left\{\hat{y}_{t+1}\right\} \\ \mathbb{E}_{j,t}^{SL}\left\{\hat{\pi}_{t+1}\right\} \end{bmatrix} = \begin{bmatrix} a_{j,t}^y \\ a_{j,t}^\pi \end{bmatrix}. \tag{5}$$

These forecast values have an appealing interpretation. In the absence of shocks, they correspond to long-run output and inflation gap forecasts. In the presence of *i.i.d.* shocks, they correspond to average output gap and inflation-gap forecasts. Under either of those interpretations, the forecasts $[a_{j,t}^y \ a_{j,t}^\pi]'$ represent agents' beliefs about the steady-state values of the inflation and output gaps, which allows us to intuitively measure expectation (un)anchoring using the distance of each variable to its respective target (*i.e.* zero).[8] Heterogeneous coefficients in $[a^y_{j,t}\ a^\pi_{j,t}]'$ capture the disagreement among forecasters observed in survey data. In particular, dispersed coefficients $\{a^\pi_{j,t}\}$ can be interpreted as disagreement about the CB's target.

Under learning, the model is solved sequentially so as to obtain a temporary equilibrium in each period, which makes it straightforward to account for the non-linearity induced by the ELB. Figure 4 summarizes the sequence of events within a period under SL. Let us now detail each step. The SL model utilizes two operators.

**Mutation.** The first operator is an innovation process, or mutation, that allows for a constant exploration of the state space outside the existing population of forecasts. Agents may experience idiosyncratic shocks, which we interpret as them receiving a new piece of information. In each period, each agent may receive news about inflation and output gaps at exogenous rates of respectively $\mu_\pi$ and $\mu_\pi$. To be more precise, each agent's forecast $m^x_{j,t}$ of any variable $x = \{y, \pi\}$ is updated at the beginning of each period as follows:

$$m^x_{j,t} = a^x_{j,t-1} + \mathbb{1}_{z^x_{j,t} \leq \mu_x}\, \iota^x_{j,t} \tag{6}$$

with $z^x_{j,t} \sim \mathcal{U}(0,1)$ a random draw from a uniform distribution with support $[0,1]$ and $\iota^x_{j,t} \sim \mathcal{N}(0, \xi^2_x)$ an idiosyncratic Gaussian random draw representing the news. Note that, the larger the parameters $\xi_x$, the wider the neighborhood to be explored around the existing forecasts, or the more disagreement there is between agents' idiosyncratic information shocks.

After the mutation process has determined a population of potential forecasts, each

[8] In the rest of the paper, we denote by $\Omega$ such an indicator of expectation anchoring. Specifically, we use the average squared distance of individual expectations to zero: $\Omega^{\mathbb{E}\pi}_t = \frac{1}{N}\sum_{j=1}^N \mathbb{E}^{SL}_{j,t}\{\widehat{\pi}_{t+1}\}^2$ and $\Omega^{\mathbb{E}y}_t = \frac{1}{N}\sum_{j=1}^N \mathbb{E}^{SL}_{j,t}\{\widehat{y}_{t+1}\}^2$. The lower those values, the stronger the anchoring of expectations.

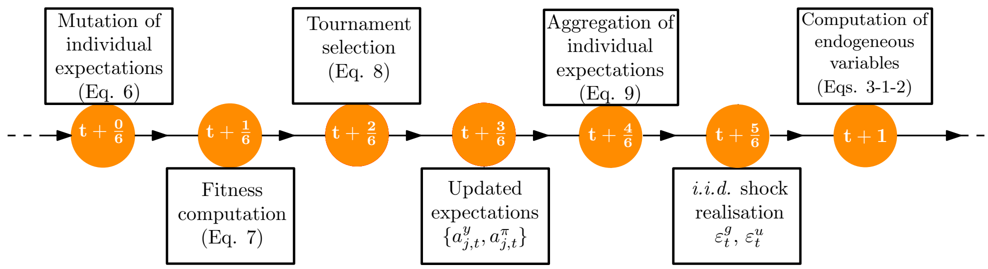


Figure 4: Intra-period timing of events in the model under SL

updated forecast $\{m^x_{j,t}\}$ enters the tournament phase, which determines the forecasts that survive evolutionary selection and actually get used by agents.

**Tournament and computation of forecasting performances.** This second operator, the tournament, is the selection mechanism of the learning process; and allows better-performing forecasts to spread among the population at the expense of worse-performing ones. Forecast performance is evaluated using forecast errors over the whole past history of the economy.

A 'fitness' value $a^x_{j,t}$, $x = \{y, \pi\}$ is assigned to each forecast a of each agent $j$. This is computed as follow

$$F^x_{j,t} = -\sum_{\tau=0}^{t} \rho^\tau_x (\widehat{x}_{t-1-\tau} - m^x_{j,t})^2. \tag{7}$$

The terms $\widehat{y}_{t-1-\tau} - m^y_{j,t}$ and $\widehat{\pi}_{t-1-\tau} - m^\pi_{j,t}$ correspond, respectively, to the output and inflation gap forecast errors that agent $j$ would have made in period $t-\tau-1$, had she used her current forecasts $m^y_{j,t}$ and $m^\pi_{j,t}$ to predict the output and inflation gaps in period $t-\tau$. The smaller the forecast errors, the higher the fitness.

Parameter $\rho_x \in [0,1]$ (for $x = y, \pi$) represents the collective memory or experience of the population. In the nested case where $\rho_x = 0$, the fitness of each forecast is completely determined by the forecast error on the most recent observable data. For any $0 < \rho_x \leq 1$, all past forecast errors impact the fitness but with exponentially declining weights while, for $\rho_x = 1$, all past errors have an equal weight in the computation of the fitness. This memory parameter allows the agents to discriminate between a one-time lucky draw and persistently good forecasting performances.

In the tournament, agents are randomly paired (the number of agents is conveniently chosen even), their fitness with respect to inflation and output gap forecasts are each compared and the agent with the lowest fitness copies the forecast of their counterpart. There are two separate tournaments: one for inflation gap forecasts $\{a^{\pi}_{j,t}\}_{j\in J}$ and one for output gap forecasts $\{a^{y}_{j,t}\}_{j\in J}$.[9] Formally, for each pair of agents $(k,l) \in J$, $k \neq l$, with a corresponding pair of forecasts $(m^x_{k,t}, m^x_{l,t})$ for each forecast variable $x \in \{\pi, y\}$, the tournament leads to an imitation of the more successful forecasts of the pair as follows:

$$(a^x_{k,t}, a^x_{l,t}) = \mathbb{1}_{F^x_{k,t} > F^x_{l,t}}(m^x_{k,t}, m^x_{k,t}) + \mathbb{1}_{F^x_{k,t} \leq F^x_{l,t}}(m^x_{l,t}, m^x_{l,t}), \qquad \text{for } x \in \{\pi, y\}. \tag{8}$$

The tournament occurs after the mutation operator in order to screen out poorly performing forecast candidates stemming from mutations. This allows the model to be less sensitive to the parameter values tuning the mutation. Indeed, if mutation were to take place after the tournament selection, all newly created forecasts would determine aggregate expectations without consideration of their performances. This way, the mutation process can be more frequent and of wider amplitude so as to allow agents to adapt more quickly

[9] This assumption will turn out useful in the empirical exercise below while not being restrictive: Arifovic et al. (2013) show that most SL dynamics are robust to a single tournament but harder to calibrate due to the systematic differential of error sizes between variables.

to new macroeconomic conditions, while limiting the amount of noise introduced by the SL algorithm.

**Aggregation of individual forecasts.** Following Arifovic et al. (2013, 2018), individual expectations (5) are aggregated using the arithmetic mean as:

$$\mathbb{E}_t^{SL}\widehat{z}_{t+1} = \frac{1}{N}\sum_{j=1}^{N}\mathbb{E}_{j,t}^{SL}\widehat{z}_{t+1}. \tag{9}$$

Appendix A.1 shows that this aggregation rule is consistent with the micro-foundations of the macroeconomic model. Note that under this aggregation procedure, agents have the same relative weight in expectations formation, thus one agent cannot influence market expectations when the number of agents $N$ is large enough.

**Computation of the endogenous variables.** Given the aggregate expectations $\mathbb{E}_t^{SL}\widehat{z}_{t+1}$ and the realization of the shocks, the piece-wise linear Taylor rule (3) sets the nominal interest rate: if the shadow rate is negative, the nominal interest rate is set to zero. Given the nominal rate, the expectations and the shock $g$, the IS curve (1) then determines the output gap and finally, the Phillips curve (2) determines the inflation gap given inflation expectations, the output gap and the shock $u$.

**Simulation protocol.** We study the dynamics of the model using numerical simulations. Throughout the rest of the paper, we proceed as described in Arifovic et al. (2013, 2018). We generate a history of 100 periods along the law of motion of the economy around the target (see Eq. (4)) and introduce a population of SL agents in $t = 100$. Their initial forecasts are drawn from the same support as the one used in the mutation process, *i.e.* from a normal distribution with standard deviation $\xi^x$, $x = \pi, y$. The first 100 periods are used to provide the agents with a history of past inflation and output gaps in order to compute the fitness of their newly introduced forecasts. In the simulation exercises in the next section, we vary the initial average of the normal distribution to tune the degree of pessimism in the economy. The further below zero the initial average forecasts are, the more pessimistic views the agents hold about future inflation and output gaps.

**Summary of the model under SL.** Combining all the equations from the model under SL, the equilibrium conditions are given by:

$$m^y_{j,t} = a^y_{j,t-1} + \mathbb{1}_{z^y_{jt} \leq \mu_y} \iota^y_{j,t} \xi_y \tag{10}$$

$$m^\pi_{j,t} = a^\pi_{j,t-1} + \mathbb{1}_{z^\pi_{jt} \leq \mu_\pi} \iota^\pi_{j,t} \xi_\pi \tag{11}$$

$$F^y_{j,t} = -\sum_{\tau=0}^{t} \rho_y^\tau (\widehat{y}_{t-1-\tau} - m^y_{j,t})^2 \tag{12}$$

$$F^\pi_{j,t} = -\sum_{\tau=0}^{t} \rho_\pi^\tau (\widehat{\pi}_{t-1-\tau} - m^\pi_{j,t})^2 \tag{13}$$

$$(a^y_{k,t}, a^y_{l,t}) = \mathbb{1}_{F^y_{k,t} > F^y_{k,t}} (m^y_{k,t}, m^y_{k,t}) + \mathbb{1}_{F^y_{k,t} \leq F^y_{l,t}} (m^y_{l,t}, m^y_{l,t}) \qquad \text{for } k \neq l \in J \tag{14}$$

$$(a^\pi_{k,t}, a^\pi_{l,t}) = \mathbb{1}_{F^\pi_{k,t} > F^\pi_{k,t}} (m^\pi_{k,t}, m^\pi_{k,t}) + \mathbb{1}_{F^\pi_{k,t} \leq F^\pi_{l,t}} (m^\pi_{l,t}, m^\pi_{l,t}) \qquad \text{for } k \neq l \in J \tag{15}$$

$$\mathbb{E}_t \left\{\hat{\pi}_{t+1}\right\} = \frac{1}{N} \sum_{j=1}^{N} a^\pi_{j,t} \tag{16}$$

$$\mathbb{E}_t \left\{\hat{y}_{t+1}\right\} = \frac{1}{N} \sum_{j=1}^{N} a^y_{j,t} \tag{17}$$

$$\hat{y}_t = \mathbb{E}_t \left\{\hat{y}_{t+1}\right\} - \frac{1}{\sigma} \left(\hat{\imath}_t - \mathbb{E}_t \left\{\hat{\pi}_{t+1}\right\}\right) + \hat{g}_t \tag{18}$$

$$\hat{\pi}_t = \kappa \hat{y}_t + \beta \mathbb{E}_t \hat{\pi}_{t+1} + \hat{u}_t \tag{19}$$

$$\hat{\imath}_t = \max\left(\phi_\pi \mathbb{E}_t \left\{\hat{\pi}_{t+1}\right\} + \phi_y \mathbb{E}_t \left\{\hat{y}_{t+1}\right\}, -\bar{r}\right) \tag{20}$$

Note that the SL model includes $6N + 2$ more equations than the RE model needs.

Finally, it is important to recognize that the RE representative-agent benchmark is nested in our heterogeneous-agent model: as soon as the inflation and output gap expectations of all agents are initialized at the targeted values and mutation is switched off (*i.e.* $\xi_y, \xi_\pi = 0$), the dynamics boil down to the RE benchmark. Under SL, our model involves a few parameters, namely the probabilities of mutation, the sizes of those mutations and the memory of the fitness function. We now detail how we match those parameter values.

# 3 Moment Matching

We will now jointly match the learning parameters and the structural parameters of the model. We first describe the construction of the dataset, then discuss our methodology and, finally, present the results.

## 3.1 Dataset

US macroeconomic time series for output, price index and nominal rates are taken from the FRED database. Forecast data come from the SPF of the Federal Reserve of Philadelphia. This choice is common in the related studies, as it is argued that these data provide a good approximation of the private sector expectations that are involved in the NK model (Del Negro & Eusepi 2011). SPF data span the period from 1968 to 2018 on a quarterly basis. To make the dataset stationary, we divide output by both the working age population and the price index. In order to obtain a measurement of the output gap, we compute the percentage deviations of the resulting output time series from its linear trend. The inflation rate is measured by the growth rate of the GDP deflator.

As heterogeneity in expectations is a key feature of real-world expectations as well as an essential ingredient of the dynamics under SL, we construct an empirical measure of that heterogeneity in the survey data. We use the cross-sectional dispersion of the individual forecasts, measured by the standard deviation of the individual forecasts among all participants in each period, to obtain a time series of forecasts' heterogeneity.

### 3.2 Methodology

With these data at hand, we proceed by matching the statistics from empirical moments with their simulated counterparts under SL. We discuss the technical details in Appendix B. In short, we use the simulated moments method (SMM), which provides a rigorous basis for evaluating whether the model is able to replicate salient business-cycle properties. The SL algorithm entails a substantial additional computational burden compared to a RE model. In particular, the SL algorithm brings an additional non-linearity into the piecewise-linear model and an additional source of stochasticity next to the fundamental shocks.[10] To circumvent these difficulties and reduce the computing time, we choose to include prior informations in the same spirit as Ruge-Murcia (2012). This moment-matching strategy, referred to as a quasi-Bayesian method, consists of a mixed estimation approach involving prior information, that aims to avoid the exploration of some parameter spaces that are economically irrelevant. With this method, priors are treated as additional moments to match in the objective function.

Hence, we first reduce the number of dimensions of the matching problem and calibrate

[10]Due to the non-linearity introduced by the ELB, we may not apply the Kalman filter and would need to use a non-linear filter to estimate the model with Bayesian full-information techniques. Given that this paper is the first attempt to bring such a heterogeneous-expectations model to the data, we encountered additional difficulties in estimating the SL model with an SMM (see Appendix B). Hence, we have left Bayesian estimation as a possibility for future research.

some of the parameters, namely the monetary policy and the preference parameters, as is standard in the related macroeconomic literature, and the number of agents (see Table 1).

We are left with four structural parameters from the NK model, namely the standard error of the fundamental shocks $\sigma_g$ and $\sigma_u$, the slope of the NK Phillips curve (parameter $\kappa$) and the natural rate $\bar{r}$. As we have calibrated the value of the discount factor $\beta$ (see Table 1), we optimize the value of the inflation trend over the period considered, which uniquely determines the value of $\bar{r}$.[11] As for the SL parameters, we need not estimate common values for the inflation and the output-gap expectation processes; as the two tournaments are separated and the two time series are likely to behave differently and exhibit different properties, both in reality and in the model. For instance, measuring inflation- and output gap-specific memory parameters $\rho_\pi$ and $\rho_y$ may reveal the fact that agents can learn that one variable may display more persistence than the other. Hence, in order to match moments, we use six learning parameters. Considering $x = \{\pi, y\}$, we optimize the mutation sizes and frequencies $\xi_x$ and $\mu_x$ as well as the memory -*i.e.* discounting rate - of the fitness measures $\rho_x$.

We now discuss the mapping between these parameters and the empirical moments to match. First, the standard deviations of the shocks $\sigma_g$ and $\sigma_u$ naturally capture the empirical volatility of output and inflation. Second, the inflation trend $\bar{\pi}$ aims to match the ELB frequency. To see why, recall that a higher natural rate $\bar{r}$ mechanically decreases the probability of hitting the ELB, as the latter is defined as $\widehat{\imath}_t = -\bar{r}$, which is strictly decreasing in the value of the inflation target. Finally, the slope of the Phillips curve $\kappa$ determines the correlation between the output and inflation gaps per Eq. (2).

As for the SL parameters, the memory parameters of the fitness function $\rho_y$ and $\rho_\pi$ tune

[11] $\bar{r}$ is associated with the inflation target per Eq. (3) but no such target existed in the US for most of the time period considered. Therefore, we use an inflation trend over that period to match the moments.

| | | Value | Source |
|---|---|---|---|
| $\sigma$ | intertemporal elasticity of substitution | 1 | Galí (2015) |
| $\varphi$ | Frish labor elasticity | 1 | Galí (2015) |
| $\phi^{\pi}$ | policy stance on inflation gap | 1.50 | Galí (2015) |
| $\phi^{y}$ | policy stance on output gap | 0.125 | Galí (2015) |
| $\beta$ | discount factor | 0.995 | Jarociński & Maćkowiak (2018) |
| $N$ | number of agents | 300 | Arifovic et al. (2018) |

Table 1: Calibrated parameters (quarterly basis)

the sluggishness of the expectations because they determine the weights on recent versus past forecast errors in the computation of the forecasting performances. The higher $\rho_y$ and $\rho_\pi$, the longer the memory of the agents, the less reactive the learning process to recent errors and the more sluggish the expectations. As sluggishness in expectations is the only source of persistence in the model once we consider *i.i.d.* shocks, parameters $\rho_y$ and $\rho_\pi$ are matched with the autocorrelation of the output and inflation gaps, respectively.

The remaining four learning parameters control the mutation processes that are the source of the pervasive heterogeneity in expectations in the SL model. We understandably use these parameters to match four moments characterizing heterogeneity in the SPF data: the average dispersion of the output and the inflation gap forecasts over the time period considered - denoted by $\Delta^{\mathbb{E}y}$ and $\Delta^{\mathbb{E}\pi}$ respectively - and their first-order autocorrelations, denoted by $\rho(\Delta_t^{\mathbb{E}y}, \Delta_{t-1}^{\mathbb{E}y})$ and $\rho(\Delta_t^{\mathbb{E}\pi}, \Delta_{t-1}^{\mathbb{E}\pi})$. In line with intuition, sensitivity analyses of the objective function of the matching problem with respect to these learning parameters have reported the following associations: the mutation sizes $\xi_y$ and $\xi_\pi$ capture a substantial share of the empirical dispersion of output and inflation gap forecasts, while the mutation frequencies $\mu_y$ and $\mu_\pi$ match most of their autocorrelation.

Finally, we impose prior restrictions on the optimized parameters and treat them as additional moments in the objective function. These restrictions allow us to incorporate additional theoretical or empirical information on parameters values in the form of priors that are combined with the data into a penalized statistical objective. The details are left for Appendix B. The priors for the structural NK parameters are taken from the literature on Bayesian estimation of DSGE models (Smets & Wouters 2007) and we choose priors for the learning parameters that are in line with the values used in the SL literature, such as Arifovic et al. (2013) (see Table 3).

## 3.3 Results from moment matching

Table 2 reports the matched moments and their empirical counterparts (in p.p.). Table 3 gives the corresponding optimized values of the parameters.

It is first striking to see that the simple two-dimensional model accounts for a substantial share of all ten moments. For half of them, the simulated moments even fall within the confidence interval of their empirical counterparts, which means that our model replicates those moments fully. We succeed in capturing not all - but a non-negligible part - of the persistence in macroeconomic variables with a model that employs only white-noise shocks.[12] We shed further light on the source of that persistence in Section 4.1. Nonetheless, at that stage, we can state that learning acts as an endogenous propagation mechanism that amplifies the effects of *i.i.d.* shocks. The learning account for 22% of the empirical output gap persistence and even 63% of the inflation persistence found in the data.

Furthermore, all simulated correlations are of the same sign as their observed counterparts. Remarkably, our model succeeds in producing positive autocorrelation in forecast

[12] Matching all the persistence would not be a realistic or desirable objective: it is unlikely that all macroeconomic persistence stems from learning in expectations, and our model ignores all other fundamental sources of persistence in the economy. We rather provide a measure of the share of the persistence that could be attributed to learning.

| Moment name | | Simulated | Empirical *[0.005;0.995]* |
|---|---|---|---|
| $\sigma(\hat{y}_t)$ | output gap sd. | 4.39 | 4.38 *[3.97;4.83]* |
| $\rho(\hat{y}_t, \hat{y}_{t-1})$ | output gap autocorr. | 0.22 | 0.98 *[0.98;0.99]* |
| $\sigma(\hat{\pi}_t)$ | inflation gap sd. | 0.66 | 0.60 *[0.54;0.66]* |
| $\rho(\hat{\pi}_t, \hat{\pi}_{t-1})$ | inflation gap autocorr. | 0.56 | 0.90 *[0.87;0.92]* |
| $\rho(\hat{\pi}_t, \hat{y}_t)$ | inflation-output correlation | 0.097 | 0.08 *[-0.07;0.21]* |
| $\Delta^{\mathbb{E}y}$ | av. forecast dispersion of output gap | 0.4 | 0.36 *[0.31;0.41]* |
| $\Delta^{\mathbb{E}\pi}$ | av. forecast dispersion of inflation gap | 0.20 | 0.25 *[0.22;0.28]* |
| $\rho(\Delta_t^{\mathbb{E}y}, \Delta_{t-1}^{\mathbb{E}y})$ | autocorr. of forecast disp. of output gap | 0.63 | 0.76 *[0.70;0.82]* |
| $\rho(\Delta_t^{\mathbb{E}\pi}, \Delta_{t-1}^{\mathbb{E}\pi})$ | autocorr. of forecast disp. of inflation gap | 0.4 | 0.64 *[0.55;0.72]* |
| $P[\hat{\imath}_t > -\bar{r}]$ | probability not at the ELB | 0.83 | 0.86 *[0.81;0.91]* |
| Objective function value | | 0.85 | - |

Table 2: Comparison of the (matched) theoretical moments with their observable counterparts

| | | | Prior Distributions | | | Posterior Results | |
|---|---|---|---|---|---|---|---|
| | | | Shape | Mean | STD | Mean | STD |
| $\sigma_g$ | - | real shock std | Invgamma | .1 | 5 | 3.8551 | 0.0774 |
| $\sigma_u$ | - | cost-push shock std | Invgamma | .1 | 5 | 0.4232 | 0.0064 |
| $100\pi^T$ | - | quarterly inflation trend | Beta | .62 | .1 | 0.829 | 0.0300 |
| $\kappa$ | - | Phillips curve slope | Beta | .05 | .1 | 0.0095 | 0.0012 |
| $\mu_y$ | - | mutation rate for $\mathbb{E}y$ | Beta | .25 | .1 | 0.2467 | 0.0035 |
| $\mu_\pi$ | - | mutation rate for $\mathbb{E}\pi$ | Beta | .25 | .1 | 0.2748 | 0.0036 |
| $\xi_y$ | - | mutation std. for $\mathbb{E}y$ | Invgamma | .1 | 2 | 0.8547 | 0.1300 |
| $\xi_\pi$ | - | mutation std. for $\mathbb{E}\pi$ | Invgamma | .1 | 2 | 0.7406 | 0.1000 |
| $\rho_y$ | - | fitness decay rate for $\mathbb{E}y$ | Beta | .5 | .2 | 0.8301 | 0.0203 |
| $\rho_\pi$ | - | fitness decay rate for $\mathbb{E}\pi$ | Beta | .5 | .2 | 0.5465 | 0.0081 |

Notes: The values of the standard deviation of the optimized parameters are computed using a numerical approximation of a sparse matrix representation of the Hessian matrix.

Table 3: Optimized parameters using the simulated method of moments matching the SPF data (1968–2018)

dispersion. This result is an important step forward in the modeling and estimation literature, as we show that our simple framework can address the heterogeneity in expectations that has been observed empirically but is incompatible with RE.

The model also does a good job at matching the probability that the ELB will bind on nominal interest rates despite the relatively small *i.i.d.* fundamental shocks. These ELB episodes do not stem from large exogenous shocks but are an endogenous product of the interplay between learning and those small *i.i.d.* shocks, as detailed in the next section.

All our optimized values are consistent with empirical values and usual estimates. For instance, the value obtained for (yearly) inflation trend is 3.4%, which falls nicely into the range between the average inflation rate over the sample.[13] Next, given the calibrated discount factor $\beta$, the implied value for the (yearly) natural interest rate is 5.45%, which is close to the average federal funds rate over the sample (namely 5.2%).

As for the values we found for the SL mutation parameters, we can see that they are all in line with the values usually employed in numerical simulations in the related literature (Arifovic et al. 2013). The optimized values of $\rho_y$ and $\rho_\pi$ imply that agents' memory is bounded,[14] which is highlighted by experimental evidence (Anufriev & Hommes 2012) and empirical estimates from micro data (Malmendier & Nagel 2016).

Our parsimonious model is therefore able to jointly and accurately reproduce ten salient features of macroeconomic time series and survey data – including the ELB duration and the pervasive heterogeneity in forecasts – while using plausible parameter values. Crucially, our model avoids diverging dynamics at the ELB without the need for exogenous restrictions on the model or the expectation process. In contrast to our model, Appendix C reproduces the same moment-matching exercise as above using a common alternative model of heterogeneous expectations based on endogenous switching between two rules: mean reversion

[13] Consider that the sample includes the 1970s (4.3%) and that the Fed inflation (2%) target was adopted later.

[14] If one discards observations weighted less than 1%, we have $0.83^{25} < 0.01$ and $0.54^{7} < 0.01$, which implies that agents' memory amounts to roughly 25 quarters for forecasting the output gap and 7 quarters for forecasting the inflation gap.

and potential extrapolation (Branch & McGough 2010, Hommes & Lustenhouwer 2019b). This alternative model can be described as an heuristic-switching (HSM) NK model. The well-known diverging deflationary dynamics at the ELB under this alternative prevent the calibration of the model other than under a setting with shadow-rate in which the occasionally binding constraint on the nominal rate is ignored. Yet, even within the shadow-rate setting, the adaptive-learning model fails to produce heterogeneity in inflation expectations (see table 7 in Appendix C). The better fit of the alternative model under shadow rates only comes from a closer match of the autocorrelation of the real variables with their empirical counterparts, but as explained above in Footnote 12, this is not a desirable or realistic outcome.

Finally, in line with Coibion & Gorodnichenko (2015), we assess the ability of each model of expectations to account for the properties of expectation errors in the SPF. To do so, in Table 4, we regress the *ex post* average forecast errors on the average forecast revisions in both the SL expectation data and the alternative heterogeneous-expectations model.[15] The resulting estimated coefficient maps into the degree of information rigidity, which may be then compared with its empirical counterpart in the SPF. Under RE, the estimated coefficients, denoted by $\beta$, should be non-significant. However, in the SPF the estimated coefficients are significantly positive (even higher than those based on inflation forecasts), which represents an under-reaction to news with a degree of information rigidity equal to $\frac{\beta}{1+\beta}$ (Coibion & Gorodnichenko 2015, Bordalo et al. 2020). Therefore, it is striking to see that SL expectations are also characterized by such a high degree of information rigidity. By contrast, the alternative model delivers negligible rigidity in inflation forecasts and even an

[15] All expectations are one-step-ahead in our model. Therefore, to obtain forecast revisions in the absence of different vintages, we interpret SL expectations as expectations of long-run inflation and output gaps that are revised over time.

| Estimated Model: | $z_{t+1} - E_t^{SL}(z_{t+1}) = c + \beta \left(E_t^{SL}(z_{t+1}) - E_{t-1}^{SL}(z_t)\right) + \text{error}_t$ | |
|---|---|---|
| Variable | Aggregate SL expectations | Aggregate HSM expectations |
| $z = \pi$ | 0.980*** | 0.132*** |
| | (0.044) | (0.013) |
| $z = y$ | 0.986*** | -0.094*** |
| | (0.125) | (0.012) |

Notes: *** sign. at 1%, ** sign. at 5%, * sign. at 10%. Estimated $\beta$-coefficients, with standard deviations between brackets. Under SL, regressions are performed over 10,000 MC simulations with fixed effects for each chain of shocks, hence the relatively small estimated standard deviations.

Table 4: Estimation of information rigidities in expectations

overreaction to news (*i.e.* a significantly negative coefficient) in output-gap forecasts.

These comparison exercises have demonstrated the interesting empirical properties of our SL model at the micro level . We may now proceed to the analysis of the underlying propagation mechanism in the model induced by SL.

# 4 Dynamics under social learning

This section first analyzes the stability properties of the targeted steady state under SL. To unravel the dynamics of expectations at the ELB, we analyze one transitory path to the target as an illustration. Next, we systematically compare the business cycle properties under SL and RE and assess the welfare loss entailed by heterogeneous expectations with respect to the RE representative agent benchmark.

## 4.1 Stability analysis

We will now examine here the asymptotic behavior of the model over the entire state space of the endogenous variables $(\hat{\pi}, \hat{y})$, as displayed in the introduction (see, again, Fig. 2). To this end, we employ Monte Carlo simulations. Figure 5 represents the phase diagram of the model where the average inflation gap expectation (*i.e.* the average of the $\{a_j^\pi\}$ values across agents) is given on the x-axis and the average output gap expectation (*i.e.* the average of the $\{a_j^y\}$ values) on the y-axis. The initial strategies are drawn around each point of the state space, and we repeat each initialization configuration 1,000 times with different seeds. We obtain the phase diagram by imposing a one-time expectational shock from the target to each point of the state space and then assess whether inflation and output gaps converge back on the targeted steady state (see Fig. 5a) and if so, at what speed (see Fig. 5b). The two figures show that the model either converges to the target (in gray-shaded areas) or diverges along a deflationary spiral (in white areas).

The main message from this exercise is that the basin of attraction of the target under SL is larger than the determinacy region of the targeted steady state under RE. To see this, notice that there is a considerable locus of points on the left-hand side of the stable manifold associated with the saddle point under recursive learning (red dashed line in Fig. 5), from which the model converges back to the target under SL.[16] By contrast, we know from the related literature that this manifold marks the frontier between (local) determinacy and indeterminacy under RE. It also marks the frontier between (local) E-stability and divergence under adaptive learning (see Evans et al. (2008) and Appendix A.3 for further details and references). This is because, under adaptive learning (or any form of purely backward-looking expectations), expectations become trend-extrapolating in this region of the state space. A single forecast in this region results in a negative forecast error ($\pi_{t+1} - E_t^{AL}(\pi_{t+1}) < 0$). Therefore, realized inflation and output gaps decline even further below their expected values

[16] As a consequence, in our model, the low-inflation state is unstable under SL as it belongs to the basin of attraction of the target: if expectations shift onto the low-inflation state, they will eventually converge back to the target. Hence, the stability result in Arifovic et al. (2018), that is obtained under infinite memory in the fitness function, does not generalize to our setup, in which agents discount past observations.

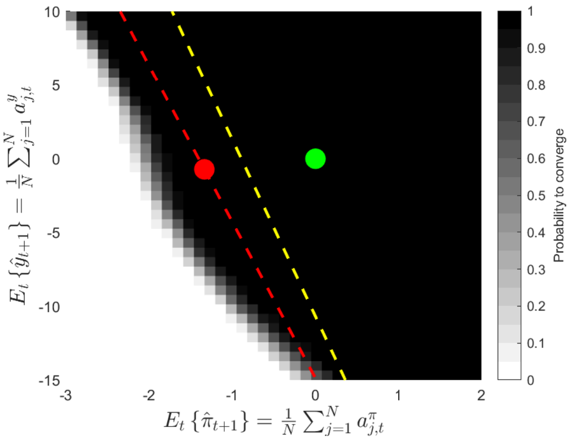


(a) Stability of the target under SL

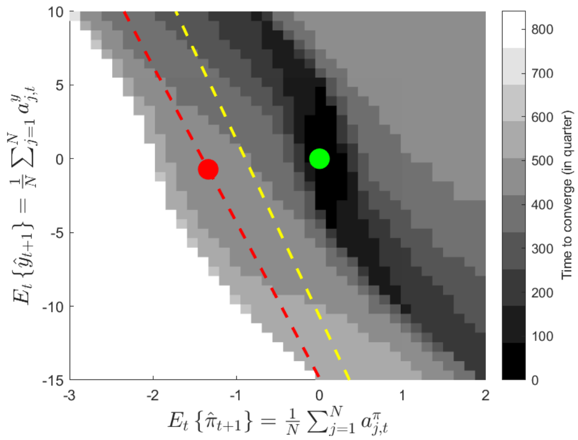


(b) Speed of convergence to the target under SL

Notes: See explanations at the end of Section 2.3. We perform 1,600,000 Monte-Carlo simulations over 1000 periods. The targeted steady state is denoted by the green dot, and the deflationary steady state by the red one. The ELB frontier (yellow dashed line) is the locus of points for which $-\overline{r} = \phi^\pi \widehat{\pi} + \phi^y \widehat{y}$: on the left-hand side, the ELB binds. The stable manifold associated with the saddle low inflation steady state (red line) is computed under recursive learning and corresponds to the stable eigenvector of $B^{elb}$: on the left-hand side, the model is indeterminate under RE and E-unstable. The empty area represents pairs of expectation values for which the model diverges along a deflationary spiral. We define convergence as $\epsilon$-convergence, i.e. inflation and output respectively enter and do not exit the neighborhood $[-\epsilon_\pi, \epsilon_\pi]$ and $[-\epsilon_y, \epsilon_y]$ with $\{\epsilon_\pi, \epsilon_y\} = \{0.1\%, 0.5\%\}$. Results are robust to tighter convergence criteria.
Left: The darker, the higher the probability to converge back to the steady state. Right: The darker, the faster the convergence back to the steady state.

Figure 5: Global dynamics under social learning

and diverge ($\pi_t - \pi_{t-1} < 0$), which causes agents to revise their expectations even further downward and eventually drives the economy along a deflationary spiral.

By contrast, under SL, individual expectations do not systematically adjust to the release of newest data points for inflation and output gaps. The best forecasters, evaluated over the (recent) history, also tend to be those who do not update their expectations much (because they are more likely to win the tournament). Even if the economy hits the ELB, the recent inflation and output experience of the agents is not consistent with a deflationary spiral (just as in real-world economies). Therefore, the strategies that would extrapolate a deflationary path (i.e. the most pessimistic ones) are not useful for predicting the recent history; as such, tend to be discarded. This is the case even after a strong pessimistic shift: as soon as some individual forecasts remain for some time above the red line in Fig. 5 (albeit below the target), these less pessimistic forecasts, which are historically more accurate, spread out and steer the economy back to the target. Conversely, our model may also lead to self-sustaining deflationary spirals when shifts in expectations are large enough to throw the entire population of strategies beyond the stable manifold. However, for this to happen, as shown by the white area in Figure 5a, the one-time shift in expectations has to be implausibly large given where the actual data lie, as depicted by Fig. 2.

Therefore, we provide an expectations model where *social interactions* aggregate *dispersed information* by relying on a *collective memory and experience*: agents evaluate their strategies with respect to a common past history. Collective memory introduces inertia process of expectation formation and updating, and favors mean-reverting forecasting strategies – even if the mean reversion may be particularly slow, as depicted in Figure 5b. Under SL, agents discard forecasting strategies that are not consistent with their common experience, a mechanism that echoes the growing literature on experience and economic decision-making (see the recent review by Malmendier & Wachter 2023).

Another noteworthy observation is given by Figure 5b. Using the same state space as Figure 5a, the figure reports the speed of convergence to the target for each pair of initial average expectations. The darker the area, the faster the convergence. It is striking to see that the closer expectations are to the targeted steady state, the faster the convergence. In general, there is a locus of points, spiraling around the target, where convergence is fast. This is consistent with the complex eigenvalues associated with that steady state.

Most interestingly, the area in the southwest of the target, beyond the stability frontier, is depicted in light gray. This means that for those severely pessimistic inflation and output gap expectations, the model under SL does converge back to the target, but does so at a particularly slow speed. This area is beyond the ELB frontier (yellow dashed line), which indicates that the ELB is binding yet the model does not diverge along a depressive downward spiral.

These observations show that our model can produce persistent but non-diverging episodes at the ELB, and heterogeneity in expectations plays an essential role in generating those dynamics. To shed more light on these dynamics, we now focus on a single expectational shock and study how it propagates in the model.

## 4.2 Illustration of persistent dynamics at the ELB

Fig. 6 illustrates the persistent dynamics at the ELB by plotting the path from one particular point of the state space back to the target.[17]

Such a shock produces a *prolonged* depressive episode at the ELB: inflation and interest rates exhibit considerable persistence below their respective targets while the output gap recovers faster, and even temporarily overshoots the steady state. These dynamics under SL are empirically much closer to the recent economic history discussed in the introduction than the excess volatility in the indeterminacy region under RE or the diverging deflationary paths under adaptive learning.

Let us now unravel the underlying forces at play under SL that deliver these empirically appealing dynamics. The initial deviations from the steady state are triggered by the pessimistic shock alone, while the resulting environment of low inflation and ELB stems entirely from the sluggish dynamics of expectations under SL and their self-fulfilling nature in the

[17] In the simulations, expectations travel to pessimistic regions of the state space as a result of the combination of SL and a series of adverse fundamental shocks. Here, we consider an arbitrary shift in expectations and use the point (0, -14)% as an example of a starting point on Fig. 5.

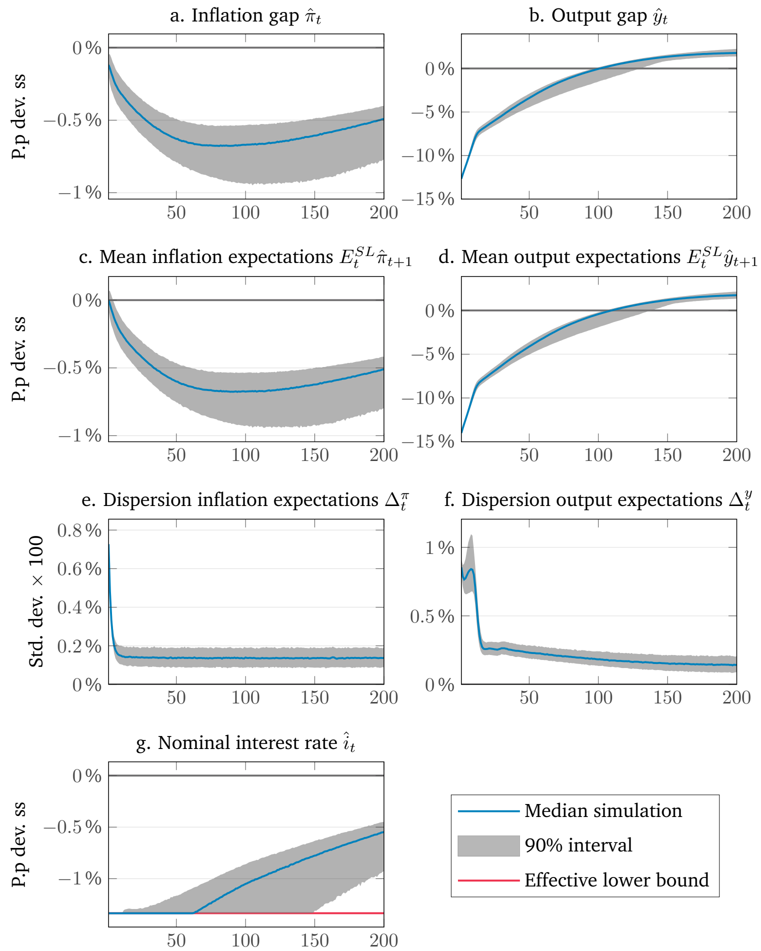


Notes: From the top left to the bottom right, transitory path of inflation gap, output gap, average inflation gap expectations, average output gap expectations, nominal interest rate, standard deviation of individual inflation expectations and standard deviation of individual output gap expectations. The solid blue line represents the median realization and the grey shaded zones are the 5% and 95% confidence intervals over 1,000 Monte Carlo simulations. All plots report the zero deviation line. The lower horizontal line in plot $g$ is the ELB.

Figure 6: Illustrative transitory path of the model after an expectation shock

NK model.

As explained in Section 4.1, right after the shock, the elimination of the most negative forecasts rules out the possibility of deflationary spirals and generates the 'missing disinflation' along the past of the bust. Per their self-fulfilling nature, below-target forecasts nurture the downturn, which triggers an accommodating response from the CB. This stimulating monetary policy has the largest impact on the output gap, which eventually turns positive.

In particular, the paths of inflation and the average inflation expectations almost perfectly overlap, which means that low inflation forecasts are almost self-fulfilling and deliver near-zero forecast errors, which allows them to diffuse among the agents. This selection mechanism explains the considerable persistence in inflation and inflation forecasts depicted in Figure 6. Inflation and inflation expectations cannot converge back to the target until the combined force of positive output gaps and low interest rates becomes strong enough to overcome the almost self-fulfilling force of low inflation expectations.[18] These dynamics generate the inflation-less recovery. This prolonged period of positive output gaps may also suggest that the economy may settle back to equilibrium only after full tapering by the CB.

Finally, it is interesting to note that our model reproduces another stylized fact discussed by Mankiw et al. (2003): a recession is associated with an increase in the dispersion of forecasts among agents – or, in other words, the level of disagreement between agents. In this simulation, the correlation between the output gap and output gap forecast dispersion is in fact significant and reaches -0.34. Indeed, Figure 6 reports how the dispersion of individual expectations spikes in the aftermath of the shock. The rise in forecast dispersion does not last; this is because the selection pressure of the SL algorithm pushes the agents to adapt to the 'new normal' in the aftermath of the shock. The level of heterogeneity between agents then returns to its long-run value, which is dictated by the size of the mutations.

[18]Admittedly, the number of periods before convergence back to the target appears implausibly large, but the model does a good job once one bears in mind that the only policy in our simple model is a Taylor rule constrained by the ELB. Our model abstracts from many empirically relevant dimensions of policy that would be likely to play a role in fostering the recovery. The simple structure of the model depicts inflation as almost entirely expectation-driven. It also ignores many other empirically relevant determinants of inflation which could also entail a quicker inflation response.

We conclude that our simple model offers a stylized representation of the observed loss of anchoring of long-run inflation expectations depicted in Figure 1. More generally, this model provides an interpretation based on expectations for the inflation dynamics in the wake of the Great Recession and the ensuing recovery as discussed in the introduction. With this model, we offer a reading of this state of affairs as the consequence of the coordination of agents' expectations on pessimistic outlooks.

From an allocation perspective, the coordination of expectations on large and persistent recessive paths pulls the economy into second-best equilibria with respect to the benchmark representative-agent model under RE.[19] Hence, SL expectations can be envisioned as a friction with respect to the RE representative-agent allocation, which may imply a substantial welfare cost, as we now demonstrate.

## 4.3 Welfare cost of social learning expectations

To evaluate this cost, we use the welfare function, which has become the main criterion, to compare alternative policy regimes. Following Woodford (2002), we consider a second-order approximation of this criterion and use the unconditional mean to express it in terms of the volatility of aggregate and idiosyncratic variables. The detailed derivations and explicit forms are deferred to Appendix A.4.

The corresponding welfare function for the average SL agent $j$ reads as:

$$E_t\left(\mathcal{W}_t\right) \simeq u_0 - u_\gamma E_j var_t\left(\hat{\gamma}_{jt}\right) - u_y var_t\left(\hat{y}_t\right) - u_\rho E_j var_t\left(\hat{\rho}_{jt}\right) - u_\pi E_j var_t\left(\hat{\pi}_{jt}\right), \tag{21}$$

[19] We refer to the RE counterpart of the NK model as the first-best equilibrium because we do not study the welfare implications of the price rigidities and imperfect competition vs. the first-best allocation under flexible prices.

where $\hat{\gamma}_{jt} = \hat{c}_{jt} - \hat{C}_t$ is the percentage deviation of the consumption of agent $j$ of aggregate consumption; $\hat{\rho}_{jt} = \hat{p}_{jt} - \hat{P}_t$ is the relative price of product $j$;[20] $u_0$ is the steady-state level of welfare; and $u_\gamma$, $u_y$, $u_\rho$ and $u_\pi$ are, the elasticities of the welfare function with respect to the variances of, in order, the idiosyncratic consumption deviation, output-gap, dispersion of relative prices and idiosyncratic inflation rates. It is straightforward to notice that macroeconomic volatility and heterogeneity among agents reduce the welfare of households.

Note that in the absence of heterogeneity across agents (such as under RE), there is no change in $\hat{\gamma}_{jt}$ and $\hat{\rho}_{jt}$ across time, thus $var\left(\hat{\gamma}_{jt}\right) = var_t\left(\hat{\rho}_{jt}\right) = 0$. In addition, inflation rates are identical across agents, which implies that $E_j var_t\left(\hat{\pi}_{jt}\right) = var\left(\hat{\pi}_t\right)$. The utility function becomes:

$$E_t\left(\mathcal{W}_t\right) \simeq u_0 - u_y var_t\left(\hat{y}_t\right) - u_\pi E_j var_t\left(\hat{\pi}_t\right). \tag{22}$$

Comparing these two allocations results in a measurement of the welfare cost of expectation miscoordination, which can be expressed in permanent consumption equivalents (Lucas 2003). Using a standard no-arbitrage condition between the SL and the RE allocations, the fraction of consumption $\lambda$ that SL households are willing to pay to live in an RE world solves the following conditions on utility streams:

$$\mathbb{E}_{j,t}^{SL}\left\{\frac{1}{N}\sum_{j=1}^{N}\sum_{t=0}^{\infty}\beta^t U\left(\left(1+\lambda\right)c_{jt}, h_{jt}\right)\right\} = \mathbb{E}_t^{RE}\left\{\sum_{t=0}^{\infty}\beta^t \mathcal{U}\left(c_t, h_t\right)\right\}. \tag{23}$$

Table 5 compares the major business cycle statistics under RE and under SL using the optimized parameters given in Table 3. This exercise allows us to disentangle the contribution

[20] Note that the consumption of agent $j$ reads as $\widehat{c}_{j,t} = \mathbb{E}^*_{j,t}\widehat{y}_{t+1} - \sigma^{-1}(\widehat{\imath}_t - \mathbb{E}^*_{j,t}\widehat{\pi}_{t+1}) + \widehat{g}_t$ and the idiosyncratic inflation rate of agent $j$ as $\widehat{\pi}_{j,t} = \beta\mathbb{E}^*_{j,t}\widehat{\pi}_{t+1} + \kappa\widehat{y}_t + \widehat{u}_t$.

| | | | Expectations scheme | |
|---|---|---|---|---|
| | | | RE | SL |
| $var\,(\hat{\pi}_t)$ | - | inflation gap variance | 0.1775 (0.002) | 0.462 (0.029) |
| $var\,(\hat{y}_t)$ | - | output gap variance | 14.816 (0.159) | 19.65 (0.644) |
| $\Delta_t^{\pi}$ | - | inflation gap forecast dispersion | – | 0.200 (0.001) |
| $\Delta_t^{y}$ | - | output gap forecast dispersion | – | 0.399 (0.002) |
| $E\,[U_t]$ | - | utility | -1.3605 (0.0001) | -1.3879 (0.002) |
| $\lambda$ | - | welfare cost | – | 0.008652 (0.0005) |
| $P[\hat{i}_t\text{=-}\bar{r}]$ | - | ELB probability | 0 (0) | 0.170 (0.026) |

Notes: Average statistics (and standard errors between brackets) over 9,400 Monte Carlo simulations of 200 periods under SL (94 series of shocks repeated 100 times) and over the same series of shocks under RE.

Table 5: Business cycle statistics and welfare under RE and SL using SMM-consistent parameters

of exogenous fluctuations in the RE-NK model from those additionally induced by SL.

Table 5 shows that SL expectations induce considerably more macroeconomic volatility than under RE, especially by inducing endogenous ELB episodes, as explained above. These self-fulfilling recessions substantially deteriorate the welfare of households in comparison to the RE benchmark. By contrast, under the assumption of *i.i.d.* shocks, the rational forecasts of inflation and output gaps boil down to their respective targeted values (see Section 2.2). Therefore, under RE, expectations remain anchored, self-fulfilling ELB episodes cannot occur and macroeconomic volatility is negligible.

The resulting cost of SL expectations with respect to RE reaches up to 0.87% of permanent consumption. This cost creates room for additional monetary policy instruments, especially communication, to enforce the additional objective of coordinating the private sector on the target.

# 5 Central bank communication

We first introduce a simple form of CB communication to develop intuition on its anchoring effect on expectations and then discuss how these insights may inform a broader range of topical communication policy debates.

## 5.1 Modeling communication under SL

We represent communication as an announcement, which we denote by $A_t^{CB}$, made by the CB at the end of any period $t$. In the model, this is an announcement about inflation in the next period $(t+1)$. We focus on inflation because it is the main objective under an inflation targeting regime.

To introduce the CB announcements into the SL algorithm, we follow Arifovic et al. (2019), albeit in a simpler game. Besides her output and inflation gap forecasts ($a_{j,t}^{y}$ and $a_{j,t}^{\pi}$), each agent $j$ now carries a probability $\psi_{j,t} \in [0,1]$ of incorporating the CB announcement into her inflation forecast in any period $t$. If she does so, her inflation forecast in $t+1$ is simply the CB announcement. Conversely, with a probability $1-\psi_{j,t}$, she ignores the announcement and sets her inflation forecast equal to her forecast $a_{j,t}^{\pi}$ as before. The determination of her output gap forecasts remains unchanged and equal to $a_{j,t}^{y}$.

Formally, in the presence of announcements, the expectation formation process of the agents given by (5) is modified as:

$$
\begin{aligned}
E_{j,t}^{SL}\{\hat{\pi}_{t+1}\} &= \begin{cases} A_t^{CB} \text{ with probability } \psi_{j,t} \\ a_{j,t}^{\pi} \text{ with probability } 1-\psi_{j,t} \end{cases} \\
E_{j,t}^{SL}\{\hat{y}_{t+1}\} &= a_{j,t}^{y}.
\end{aligned} \tag{24}
$$

The communication-augmented inflation forecast $\{(\psi_{j,t}, a^{\pi}_{j,t})\}_{j\in J}$ undergoes the same mutation and tournament processes as the output gap forecast $a^{y}_{j,t}$ (see Section 2.3).[21] The only difference from the algorithm used so far lies in the computation of the fitness of inflation forecasts, where Eq. (7) is modified as follows:

$$F^{\pi}_{j,t} = -\psi_{j,t}\sum_{\tau=0}^{t}\rho_{\pi}^{\tau}(\widehat{\pi}_{t-1-\tau} - A^{CB}_{t-\tau-1})^2 - (1-\psi_{j,t})\sum_{\tau=0}^{t}\rho_{\pi}^{\tau}(\widehat{\pi}_{t-1-\tau} - m^{\pi}_{j,t})^2,$$

where the first (resp. second) term now corresponds to the discounted sum of squared forecast errors had the agent followed (resp. ignored) the announcements of the CB.

The probabilities $\{\psi_j\}$ can be easily interpreted as the credibility of the announcements. If agents following the announcements (*i.e.*, agents with a relatively high value of $\psi_j$) have lower forecast errors than agents ignoring the announcements (*i.e.* agents with a relatively low value of $\psi_j$), then following as a strategy will spread among agents, which means that the average value of $\psi$ across agents will increase. The opposite holds if following the announcements performs more poorly than ignoring them. Thus, SL agents endogenously build trust or distrust in the communication of the CB as a function of the relative forecasting performances of each alternative. We now develop two simple examples of announcements to show how communication affects expectations.

## 5.2 Two simple communication examples

We consider the following two communication examples under SL.

[21] In the simulations below, the initial credibility $\{\psi_{j,0}\}$ is drawn from a normal distribution centered around 0.5 with a standard deviation equal to 0.25, a value that is also taken to dictate the mutation process of the probabilities $\{\psi_{j,0}\}$. Results are insensitive to alternatives.

**The CB announces the inflation target** We then have $A_t^{CB} = 0$ (as the model is written in deviations from the steady state). It should be noted that the target corresponds to the RE inflation forecasts in our simple model. The announcement of the CB is therefore consistent with the conduct of monetary policy under RE. Hence, the inflation target is redundant information to RE agents, but this piece of information may play a non-trivial role under SL.

**The CB announces its own inflation forecasts for the next period** We assume that the policy authority estimates a commonly used VAR forecasting model that is recursively updated with new observations in each period. Note that assuming VAR forecasting amounts to assuming that the CB is aware of agents being boundedly rational and, therefore, includes past realizations of the endogenous variables in its forecasting model to account for the propagation mechanism induced by learning. Indeed, such a forecasting model would be misspecified should the agents have RE and, hence, the economy evolve according to the MSV solution. In this second communication scenario, the announcement of the CB is therefore consistent with the conduct of monetary policy under SL.[22]

We now develop intuitions on how communication affects agents' expectations under SL. First, Table 6 compares the business cycle statistics of the model under RE and SL – for ease of reading, the first two columns recall the statistics in Table 5 – and under the two communication scenarios, *i.e.*, when the target and the inflation forecasts are announced.

The first three rows of Table 6 indicate that communication significantly improves macroe-

[22] The MSV solution under SL is a complicated and non-linear function of all the states in the system, including those pertaining to the SL process, and an explicit form is not available. We claim that the best the CB can do in such an environment is to estimate the law of motion of the economy with an atheoretical model, such as a VAR. We choose 8 lags, in line with the memory of the agents that is implied by the optimized value of the fitness memory on inflation (see, again, Table 3). Results are robust to more or fewer lags.

| | *RE* No communication | *SL* No communication | *SL* target | *SL* VAR(8) forecast |
|---|---|---|---|---|
| **Macroeconomic variability** | | | | |
| $var\ (\hat{\pi}_t)$ | 0.178 (*0.002*) | 0.463(*0.029*) | 0.215 (*0.010*) | 0.245 (*0.010*) |
| $var\ (\hat{y}_t)$ | 14.81 (*0.159*) | 19.65 (*0.644*) | 17.18 (*0.310*) | 17.12 (*0.304*) |
| **ELB frequency** | | | | |
| $P[\hat{i}_t\text{=-}\bar{r}]$ | 0.000 (*0.000*) | 0.170 (*0.026*) | 0.038 (*0.012*) | 0.002 (*0.001*) |
| **Expectations dispersion** | | | | |
| $\Delta_t^{\pi}$ | - | 0.200 (*0.001*) | 0.138 (*0.004*) | 0.155(*0.0045*) |
| $\Delta_t^{y}$ | - | 0.399 (*0.002*) | 0.396 (*0.002*) | 0.397 (*0.002*) |
| **Expectations anchoring** | | | | |
| $\Omega_t^{\pi}$ | 0.000 (*0.000*) | 0.887 (*0.107*) | 0.108 (*0.020*) | 0.176 (*0.003*) |
| $\Omega_t^{y}$ | 0.000 (*0.000*) | 10.03 (*1.220*) | 6.433 (*0.336*) | 6.310 (*0.359*) |
| **Welfare cost** | | | | |
| $E\left[U_t\right]$ | -1.3605 (*0.0001*) | -1.3879 (*0.001*) | -1.3732 (*0.001*) | -1.3767 (*0.001*) |
| $\lambda$ | 0.00 | 0.008652 (0.0005) | 0.003957 (0.0002) | 0.005061 (0.0002) |

Notes: See Table 5.

Table 6: Business cycle statistics under RE, under SL and with CB communication about the inflation target and the inflation forecasts

conomic stabilization with respect to the baseline SL model: the volatility of inflation decreases by more than 50% and the risk of ELB episodes drops considerably. A look at the next four lines of Table 6 reveals that not only are expectations better coordinated (*i.e.* disagreement between agents is reduced) in the presence than in the absence of communication, but coordination occurs around the CB objectives (*i.e.* expectations are better anchored at the target).

Hence, we first conclude that in our model, CB communication acts as an anchor for heterogeneous expectations and, by improving their coordination, communication contributes to macroeconomic stabilization. This effect translates into a narrower, yet positive, welfare

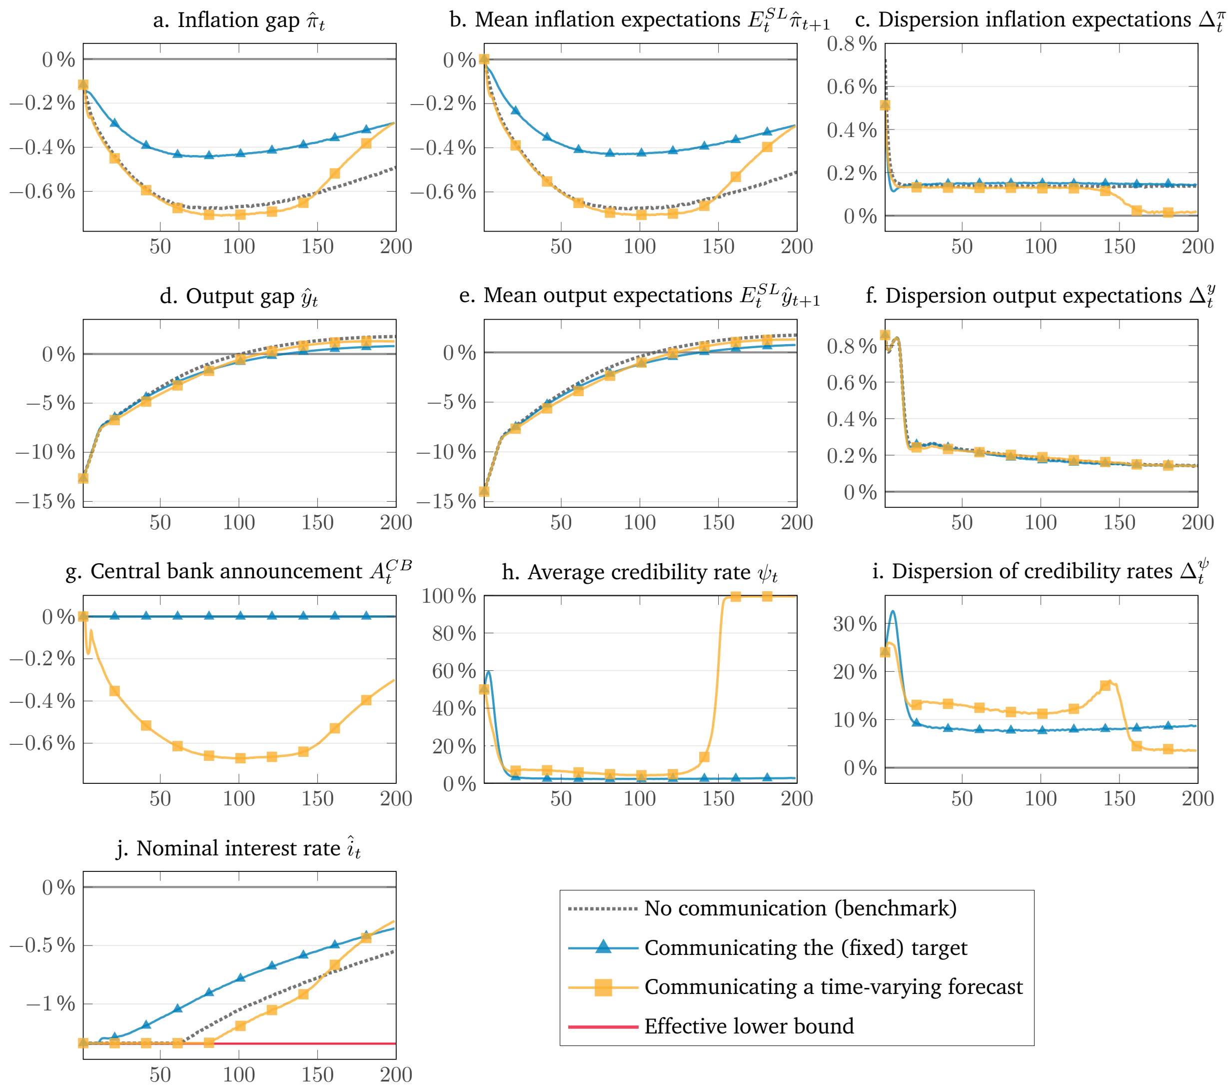


Notes: See Figure 6. The median realizations over 1,000 Monte Carlo simulations are reported. The grey line corresponds to the baseline SL model without communication. The blue and orange lines report the scenarios with communication about the CB forecasts (with blue triangles) and the target (with orange squares).

Figure 7: Illustrative transitory path of the model to a one-period -14% output gap expectation shock under various communication scenarios

gap with respect to the RE representative agent benchmark.

Next, we consider the same illustrative transitory path as in Section 4.2 with communication; see Figure 7. In the wake of the shock, both communication scenarios result in a loss of credibility. As a consequence, both types of announcement temporarily lose their anchoring power on agents' inflation expectations; see Figure 7h, where credibility invariably drops towards zero right after the shock. When announcing the target, this credibility loss stems from the actual realizations of inflation drifting away from the target. When announcing forecasts, the credibility loss results from the inaccuracy of the announced forecasts, as the pessimistic shock is unexpected – to see that, look at the discrepancy between the plunging inflation and the near-target announcements immediately after the shock (Fig. 7a vs. 7g). In both cases, the forecasting performances of the followers deteriorate and a large fraction of the agents stop following the CB's announcements.

This credibility loss leads us to the second conclusion: in our model, agents need to 'see it to believe it'. In other words, if the CB's announcements are decoupled from the actual inflation dynamics, they lose their anchoring power on expectations.

Next, as time goes forward, the CB, by updating its model, provides more accurate forecasts and regains credibility. To see that, notice the similarity between the announcements and actual inflation some periods after the shock. At the same time, this coordination on the forecast announcements leads to a reduction in expectation heterogeneity – to see this, notice the drop in inflation forecast dispersion (Fig. 7c) as credibility increases (Fig. 7h). By contrast, if just announcing the target, the CB only regains its credibility once inflation has converged back to the target, which may take a considerable amount of time, as discussed in Section 4.

Yet announcing forecasts is not a panacea: doing so also accentuates the downturn. Indeed, inflation dives deeper, the ELB binds for a longer period (Fig. 7j) and output overshoots further (Figs. 7d-7e) than when the CB announces its target. This observation illustrates an important pitfall of communication: by extrapolating the bust, the announced forecasts may turn self-defeating due to the self-fulfilling nature of expectations, and contribute to driving expectations away from the target. This striking effect is illustrated in the three graphs of the average inflation forecasts (Fig. 7b), the CB forecast announcements (Fig. 7g) and the actual inflation (Fig. 7a), all of which almost overlap.

## 5.3 A broader policy perspective

From these two simple communication examples, we can inform a broader range of monetary policy issues. For instance, our framework can inform the 'forward-guidance puzzle' (Carlstrom et al. 2012): under RE, any CB announcement about the future is immediately incorporated into agents' expectations and optimal decisions and triggers dramatic effects right from the time of the announcement. Yet, empirical evidence contradicts such a strong effect; see, *inter alia*, (Del Negro et al. 2012, Campbell et al. 2016). Based on our SL model, imperfect credibility may play a central role in explaining this puzzle.[23] If agents need to 'see it first to believe it', announcements that are at odds with the actual inflation dynamics have a milder effect on expectations and hence actual decisions than under the full credibility assumption underlying RE.

Another example of a topical policy debate that can be informed by our work is the discussion about history-dependent rules, such as average-inflation targeting or price-level targeting, which imply time-varying inflation targets. Per the self-fulfilling nature of inflation expectations, inflation expectations need to be consistent with the new policy rule; in other words, the desired future inflation rates need to be understood by the public and the ability of the CB to deliver them needs to be credible. Those considerations reinforce the rationale

[23]Other solutions in the literature rely on weakening the effect of expected real interest rates on consumption by adding frictions such as liquidity constraints, limited asset-market participation or habit formation Del Negro et al. (2012). Related contributions explain the puzzle by weakening the expectation channel if agents use k-level reasoning (Farhi & Werning 2019) or pay limited attention (Gabaix 2020).

for intensifying the CBs' efforts to communicate with the public.

Finally, announcing the CB's forecasts can be envisioned within the inflation-forecast targeting (IFT) framework. IFT is based on the principle that, given a long-term inflation objective, the CB's own inflation forecasts act as time-varying intermediate targets because such a forecast path embodies all the relevant information available to the policy makers. It has been conceived as a way to circumvent the rigidity of a purely rule-based reaction function while avoiding any expectations drifts that may result from a discretionary approach (Woodford 2007).

Our communication exercise shows how IFT allows the CB to coordinate inflation expectations despite the indeterminacy generated by the neutralization of the interest rate feedback at the ELB. Communication may also be particularly relevant in a high-inflation environment. IFT became particularly useful in the case of transition economies adopting an inflation-targeting regime; see, e.g., Clinton et al. (2017) for the case of Czech Republic. In this case, the CB aims to bring inflation to the newly announced target and anchor expectations there. It does so by announcing inflation forecasts that gradually converge to the target in an attempt to coordinate expectations on these forecasts and gradually steer inflation and inflation expectations towards the target. In the aftermath of the COVID-19 pandemic – characterized by elevated, broad-based and persistent inflation – such a communication is certainly part of the toolbox of CBs to reduce disagreement between agents and steer their inflation expectations back on target.

# 6 Conclusion

This paper develops a micro-founded model that features expectations-driven business cycles. The key mechanism works through heterogeneous expectations that may lose their anchoring to the target and persistently coordinate on below-target paths, which triggers prolonged ELB episodes. Heterogeneous expectations are introduced via an SL process into an otherwise standard two-equation macroeconomic model with a constrained Taylor rule. Our model nests the RE representative agent benchmark. In particular, we use white-noise fundamental shocks to isolate the contribution of non-rational expectations in the formation of the business cycle.

Our first contribution is to bring such a model to the data and jointly optimize its fundamental and learning parameters to match moments from both US inflation and output gaps and the SPF. Our parsimonious model is able to account for ten stylized facts, including properties related to heterogeneity in forecasts, persistence in macroeconomic variables and the endogenous occurrence of ELB episodes.

We then analyze the dynamics of the model and show that the basin of attraction of the target under SL is larger than the determinacy region under RE. In the context of our model, ELB episodes are times in which expectations have coordinated on pessimistic outlooks following a series of adverse fundamental shocks. In another word, expectations have visited regions of that basin from where the transition back on the target does occur but at a particularly slow pace. Our second major contribution is then to provide a framework that can account for the inflation experience from the aftermath of the Great Financial Crisis to the COVID-19 pandemic that is challenging to capture in macroeconomic models. In particular, our model accounts for the 'missing disinflation' during the Great Recession, per its stable but below-target dynamics and extensive ELB episodes. It also accounts for the ‘inflation-less recovery’ resulting from the combination of unanchored inflation expectations – which put downward pressure on inflation – and the boosting effect of low interest rates on output.

Finally, we extent our model to illustrate how CB communication may influence expectations. In our model, the credibility of the announcements is not *a priori* granted but rather follows the same evolutionary process as the forecasts of the agents. From two simple examples, we show that this endogenous credibility plays a central role in reshaping expectations: in our model, agents need to ‘see it to believe it’. Moreover, pessimistic inflation forecasts may turn self-defeating, per the self-fulfilling nature of inflation expectations. From these observations, we discuss broader policy implications to shed light on recent debates. Our exercise is relevant to address issues such as the forward-guidance puzzle, the implementation of history-dependent rules or address current challenges such as volatile and elevated inflation.

Our model offers a simple framework that nevertheless opens up the possibility for analyzing a rich set of monetary policy alternatives. As for our optimization routine, it may be applied to a wide range of standard workhorse models, which could then be explored under heterogeneous expectations. Those research avenues are left for future work.

# References


Andrade, P., Gaballo, G., Mengus, E. & Mojon, B. (2019), 'Forward guidance and heterogeneous beliefs', *American Economic Journal: Macroeconomics* **11**(3), 1–29.

Angeletos, G.-M., Collard, F. & Dellas, H. (2018), 'Quantifying confidence', *Econometrica* **86**(5), 1689–1726.

Anufriev, M., Assenza, T., Hommes, C. & Massaro, D. (2013), 'Interest rate rules and macroeconomic stability under heterogeneous expectations', *Macroeconomic Dynamics* **17**(8), 1574–1604.

Anufriev, M. & Hommes, C. (2012), 'Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments', *American Economic Journal: Microeconomics* **4**(4), 35–64.

Arifovic, J., Boitnott, J. F. & Duffy, J. (2019), 'Learning correlated equilibria: An evolutionary approach', *Journal of Economic Behavior & Organization* **157**(C), 171–190.

Arifovic, J., Bullard, J. & Kostyshyna, O. (2013), 'Social learning and monetary policy rules', *The Economic Journal* **123**(567), 38–76.

Arifovic, J., Schmitt-Grohé, S. & Uribe, M. (2018), 'Learning to live in a liquidity trap', *Journal of Economic Dynamics and Control* **89**, 120–136.

Aruoba, B., Cuba-Borda, P. & Schorfheide, F. (2017), 'Macroeconomic dynamics near the ZLB: A tale of two countries', *The Review of Economic Studies* **85**(1), 87–118.

Aruoba, S. B., Mlikota, M., Schorfheide, F. & Villalvazo, S. (2022), 'Svars with occasionally-binding constraints', *Journal of Econometrics* **231**(2), 477–499.

Ascari, G. & Sbordone, A. M. (2014), 'The macroeconomics of trend inflation', *Journal of Economic Literature* **52**(3), 679–739.

Benhabib, J., Schmitt-Grohé, S. & Uribe, M. (2001), 'The perils of Taylor rules', *Journal of Economic Theory* **96**(1), 40–69.

Bordalo, P., Gennaioli, N., Ma, Y. & Shleifer, A. (2020), 'Overreaction in macroeconomic expectations', *American Economic Review* **110**(9), 2748–82.

Branch, W. A. (2004), 'The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations', *The Economic Journal* **114**(497), 592–621.

Branch, W. A. & Gasteiger, E. (2019), Endogenously (non-) ricardian beliefs, Technical report, ECON WPS.

Branch, W. A. & McGough, B. (2009), ‘A new keynesian model with heterogeneous expectations’, *Journal of Economic Dynamics and Control* **33**(5), 1036–1051.

Branch, W. A. & McGough, B. (2010), ‘Dynamic predictor selection in a new Keynesian model with heterogeneous expectations’, *Journal of Economic Dynamics and Control* **34**(8), 1492–1508.

Brock, W. A. & Hommes, C. H. (1997), ‘A rational route to randomness’, *Econometrica: Journal of the Econometric Society* pp. 1059–1095.

Bullard, J. & Mitra, K. (2002), ‘Learning about monetary policy rules’, *Journal of Monetary Economics* **49**(6), 1105–1129.

Busetti, F., Delle Monache, D., Gerali, A. & Locarno, A. (2017), Trust, but verify. de-anchoring of inflation expectations under learning and heterogeneity, Technical report, ECB working paper.

Campbell, J. R., Fisher, J. D. M., Justiniano, A. & Melosi, L. (2016), Forward guidance and macroeconomic outcomes since the financial crisis, *in* ‘NBER Macroeconomics Annual 2016’, Vol. 31, University of Chicago Press, pp. 283–357.

Carlstrom, C. T., Fuerst, T. S. & Paustian, M. (2012), ‘Inflation and output in New Keynesian models with a transient interest rate peg’, *Journal of Monetary Economics* **76**, 230 – 243.

Carroll, C. D. (2003), ‘Macroeconomic expectations of households and professional forecasters’, *Quarterly Journal of economics* **118**(1), 269–298.

Clinton, K., Hlédik, T., Holub, M. T., Laxton, M. D. & Wang, H. (2017), *Czech Magic: Implementing Inflation-Forecast Targeting at the CNB*, International Monetary Fund.

Coibion, O. & Gorodnichenko, Y. (2015), ‘Information rigidity and the expectations formation process: A simple framework and new facts’, *American Economic Review* **105**(8), 2644–78.

Coibion, O., Gorodnichenko, Y., Kumar, S. & Pedemonte, M. (2020), ‘Inflation expectations as a policy tool?’, *Journal of International Economics* **124**, 103297.

Coibion, O., Gorodnichenko, Y. & Weber, M. (2019), Monetary policy communications and their effects on household inflation expectations, Technical report, National Bureau of Economic Research.

Cornea-Madeira, A., Hommes, C. & Massaro, D. (2019), ‘Behavioral heterogeneity in us inflation dynamics’, *Journal of Business & Economic Statistics* **37**(2), 288–300.

De Grauwe, P. (2011), 'Animal spirits and monetary policy', *Economic theory* **47**(2), 423–457.

Del Negro, M. & Eusepi, S. (2011), 'Fitting observed inflation expectations', *Journal of Economic Dynamics and Control* **35**(12), 2105–2131.

Del Negro, M., Giannoni, M. & Patterson, C. (2012), The forward guidance puzzle, Staff Reports 574, Federal Reserve Bank of New York.

Di Bartolomeo, G., Di Pietro, M. & Giannini, B. (2016), 'Optimal monetary policy in a new keynesian model with heterogeneous expectations', *Journal of Economic Dynamics and Control* **73**, 373–387.

Eusepi, S. & Preston, B. (2010), 'Central bank communication and expectations stabilization', *American Economic Journal: Macroeconomics* **2**(3), 235–71.

Evans, G. W., Guse, E. & Honkapohja, S. (2008), 'Liquidity traps, learning and stagnation', *European Economic Review* **52**(8), 1438–1463.

Evans, G. W., Honkapohja, S. & Mitra, K. (2022), 'Expectations, Stagnation, and Fiscal Policy: a nonlinear analysis', *International Economic Review* **63**(3).

Farhi, E. & Werning, I. (2019), 'Monetary Policy, Bounded Rationality, and Incomplete Markets', *American Economic Review* **109**(11), 3887–3928.

Gabaix, X. (2020), 'A behavioral New Keynesian model', *American Economic Review* **110**(8), 2271–2327.

Galí, J. (2015), *Monetary policy, inflation, and the business cycle: An introduction to the New Keynesian framework and its applications*, Princeton University Press.

Gasteiger, E. (2014), 'Heterogeneous Expectations, Optimal Monetary Policy, and the Merit of Policy Inertia', *Journal of Money, Credit and Banking* **46**(7), 1535–1554.

Goy, G., Hommes, C. & Mavromatis, K. (2020), 'Forward guidance and the role of central bank credibility under heterogeneous beliefs', *Journal of Economic Behavior & Organization* .

Greenwood, J., Hercowitz, Z. & Huffman, G. W. (1988), 'Investment, capacity utilization, and the real business cycle', *The American Economic Review* pp. 402–417.

Guerrieri, L. & Iacoviello, M. (2015), 'Occbin: A toolkit for solving dynamic models with occasionally binding constraints easily', *Journal of Monetary Economics* **70**, 22–38.

Gust, C., Herbst, E., LÃşpez-Salido, D. & Smith, M. E. (2017), 'The Empirical Implications of the Interest-Rate Lower Bound', *American Economic Review* **107**(7), 1971–2006.

Hachem, K. & Wu, J. C. (2017), 'Inflation announcements and social dynamics', *Journal of Money, Credit and Banking* **49**(8), 1673–1713.

Hansen, L. P. (1982), 'Large sample properties of generalized method of moments estimators', *Econometrica* pp. 1029–1054.

Hansen, N., Müller, S. D. & Koumoutsakos, P. (2003), 'Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES)', *Evolutionary Computation* **11**(1), 1–18.

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

Hommes, C. H. & Lustenhouwer, J. (2019*a*), 'Managing unanchored, heterogeneous expectations and liquidity traps', *Journal of Economic Dynamics and Control* **101**, 1–16.

Hommes, C. & Lustenhouwer, J. (2019*b*), 'Inflation targeting and liquidity traps under endogenous credibility', *Journal of Monetary Economics* **107**, 48–62.

Hommes, C., Massaro, D. & Weber, M. (2019), 'Monetary policy under behavioral expectations: Theory and experiment', *European Economic Review* **118**, 193–212.

Jarociński, M. & Maćkowiak, B. (2018), 'Monetary-fiscal interactions and the euro area's malaise', *Journal of International Economics* **112**, 251–266.

Lansing, K. J. (2020), 'Endogenous forecast switching near the zero lower bound', *Journal of Monetary Economics* .

Lindé, J. & Trabandt, M. (2019), Resolving the Missing Deflation Puzzle, CEPR Discussion Papers 13690, C.E.P.R. Discussion Papers.

Lucas, R. E. (2003), 'Macroeconomic priorities', *American Economic Review* **93**(1), 1–14.

Malmendier, U. & Nagel, S. (2016), 'Learning from inflation experiences', *The Quarterly Journal of Economics* **131**(1), 53–87.

Malmendier, U. & Wachter, J. (2023), Memory of Past Experiences and Economic Decisions, *in* M. Kahana & A. W. (eds.), eds, 'Handbook of Human Memory', Oxford University Press.

Mankiw, N. G., Reis, R. & Wolfers, J. (2003), 'Disagreement about inflation expectations', *NBER Macroeconomics Annual* **18**, 209–248.

Massaro, D. (2013), 'Heterogeneous expectations in monetary DSGE models', *Journal of Economic Dynamics and Control* **37**(3), 680–692.

McFadden, D. (1989), 'A method of simulated moments for estimation of discrete response models without numerical integration', *Econometrica* **57**(5), 995–1026.

Milani, F. (2007), 'Expectations, learning and macroeconomic persistence', *Journal of Monetary Economics* **54**(7), 2065–2082.

Orphanides, A. & Williams, J. C. (2004), *Imperfect Knowledge, Inflation Expectations, and Monetary Policy*, University of Chicago Press, pp. 201–246.

Orphanides, A. & Williams, J. C. (2005), 'Inflation scares and forecast-based monetary policy', *Review of Economic Dynamics* **8**(2), 498–527.

Orphanides, A. & Williams, J. C. (2007), 'Robust monetary policy with imperfect knowledge', *Journal of monetary Economics* **54**(5), 1406–1435.

Ozden, T. (2021), 'Heterogeneous expectations and the business cycle at the effective lower bound'.

Ozden, T. & Wouters, R. (2021), 'Restricted Perceptions, Regime Switches and the Effective Lower Bound'.

Ruge-Murcia, F. (2012), 'Estimating nonlinear dsge models by the simulated method of moments: With an application to business cycles', *Journal of Economic Dynamics and Control* **36**(6), 914 – 938.

Schmitt-Grohé, S. & Uribe, M. (2017), 'Liquidity traps and jobless recoveries', *American Economic Journal: Macroeconomics* **9**(1), 165–204.

Slobodyan, S. & Wouters, R. (2012*a*), 'Learning in a medium-scale DSGE model with expectations based on small forecasting models', *American Economic Journal: Macroeconomics* **4**(2), 65–101.

Slobodyan, S. & Wouters, R. (2012*b*), 'Learning in an estimated medium-scale DSGE model', *Journal of Economic Dynamics and Control* **36**(1), 26–46.

Smets, F. & Wouters, R. (2007), 'Shocks and frictions in US business cycles: A Bayesian DSGE approach', *American Economic Review* **97**, 586–606.

Woodford, M. (2002), 'Inflation stabilization and welfare', *The B.E. Journal of Macroeconomics* **2**(1), 1–53.

Woodford, M. (2007), 'The case for forecast targeting as a monetary policy strategy', *Journal of Economic Perspectives* **21**(4), 3–24.

## Online Appendix

(not for publication)

# A Derivations of the model

## A.1 A micro-founded heterogenous expectations model

We develop here a standard NK model with heterogeneous expectations shaped by social interactions. The time dimension and the number of firms and households are discrete. Specifically, we assume that there are $N$ firms and households, indexed by $j = 1, \ldots, N$, that are similar (in particular in terms of preferences and technology) except when it comes to their inflation and output expectations. Each firm is permanently matched to a household $j$ which owns it, such that expectations of agent $j$ refer to the corresponding household-firm.

### A.1.1 Heterogeneous expectations under SL

Under the social-learning (SL) algorithm of Arifovic et al. (2013), the information set of the agents differs from the one prevailing under RE. Let $\mathbb{E}_t^* \{\cdot\}$, $* = \{RE, SL\}$, denote the expectation operator for RE and SL, respectively.. Let $x$ denote the variable that agents in the model forecast. Agent $j$'s expectation is given by:

$$\mathbb{E}_{j,t}^{SL} \{x_{t+1}\} = \exp(a_{j,t}^x)\mathbb{E}_t^{RE} \{x_{t+1}\}, \tag{25}$$

where $\mathbb{E}_t^{RE} \{x_{t+1}\}$ corresponds to the rational expectation of variable $x$ for $t+1$ given the information set available at time $t$ and $\exp(a_{j,t}^x)$ is the idiosyncratic information of agent $j$ about future realizations of $x$ (affected by the SL algorithm up until period $t-1$). It is assumed from the SL expectation scheme that $\exp(a_{j,t}^x)$ follows some stochastic process with mean 1.

In the absence of autocorrelation in exogenous disturbances, there is no state variable in the RE solution. As a consequence, expectations are always anchored to their steady-state values, namely $\mathbb{E}_t^{RE} \{x_{t+1}\} = x^T$, for $x = \pi, y$. Consequently, Equation 25 becomes:

$$\mathbb{E}_{j,t}^{SL} \{x_{t+1}\} = \exp(a_{j,t}^x)x^T. \tag{26}$$

Taking a first-order approximation of Equation 25 around the steady state yields the following expression:

$$\mathbb{E}_{j,t}^{SL} \{\hat{x}_{t+1}\} \simeq a_{j,t}^x, \tag{27}$$

where variables with a hat denote log-deviations from their steady-state values.

### A.1.2 Households

Our economy is populated by an infinitely-lived family composed of a discrete number $N$ of members, indexed by $j \in [1, ..., N]$. Each member decides about their consumption, labor and saving plans in order to maximize the household's welfare:

$$\frac{1}{N}\sum_{j=1}^{N}\mathbb{E}_{j,t}^{*}\sum_{\tau=0}^{\infty}\beta^{\tau}\frac{1}{1-\sigma'}\left(c_{jt+\tau}-\chi\frac{h_{jt+\tau}^{1+\varphi}}{1+\varphi}\right)^{1-\sigma'}. \tag{28}$$

The individual utility is increasing in consumption $c_{j,t}$ and decreasing in labor $h_{j,t}$, where $\sigma'$ and $\varphi$ are curvature parameters and $\beta$ is the discount factor. Note that the utility function is non-separable, based on the utility function of Greenwood et al. (1988). This specification is convenient in heterogeneous-agent models as non-separability mutes the wealth effect in the labor supply, which ensures that all members of the family supply the same amount of hours $h$. Because $N$ is large enough, each agent assumes that their decisions do not impact the aggregate variables.

Family members face an intertemporal problem: they determine the levels of consumption $c_{j,t}$, hours worked $h_{j,t}$ and real-bond holdings $b_{j,t}$ which maximize the welfare of the family under the following budget constraint which binds in every period:

$$c_{j,t}+b_{j,t}=\frac{i_{t-1}}{\pi_t}\frac{b_{j,t-1}}{\exp\left(\varsigma_g\hat{g}_t\right)}+\Pi_{j,t}+w_t h_{j,t}+T_{j,t}, \tag{29}$$

where $w_t$ is the real wage (symmetric across members as they all have the same marginal product of labor); $i_{t-1}$ the nominal interest rate payable on bond holdings $b_{j,t-1}$; $\Pi_{j,t}$ is the share of the real profits from monopolistic competition paid to member $j$; $\pi_t$ the inflation rate between periods $t-1$ and $t$; and $T_{j,t}$ the lump-sum government transfers, which may be positive or negative. Variable $\hat{g}_t$ denotes an exogenous source of aggregate fluctuations, referred to as the risk-premium shock in the Smets & Wouters (2007) model, and is affected by the elasticity parameter $\varsigma = \sigma'\frac{(1-\chi)}{\vartheta}$, which normalizes the shock in the linearized version of the aggregate-demand equation.

Agents choose their consumption and savings plans ($c_{j,t}$,$b_{j,t}$) conditional on their inflation- and output-gap expectations. Hence, heterogeneity in expectations may entail heterogeneous consumption values and wealth across agent types $j$, which poses a challenge for aggregation, in particular when saving is used in the production through capital goods. With labor as the unique input in the production function and in the absence of any borrowing constraint, the level of consumption is solely determined by the Euler equation, in which the idiosyncratic saving stock does not directly affect consumption patterns.[24] The idiosyncratic saving stocks $b_{j,t}$ may be positive if saving or negative if borrowing. Nonetheless, each family member is not allowed to run a Ponzi scheme. Therefore we do not need to track the individual $b_{j,t}$ because they do not affect consumption under the assumptions of the model. The aggregate demand for government bonds reads as follows: $B_t = \sum_{j=1}^{N} b_{j,t}$. Finally, note that as in learning models with an ELB constraint, the transversality condition may not hold locally. If the ELB binds, agents may postpone consumption, which reinforces the deflation and sends the economy into a deflationary spiral. This may happen, albeit infrequently, whenever a negative shocks on the expectations is too large, as explained in Section 4.1.

Each household $j$ solves the following problem:

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

The first-order conditions are given by:

$$w_t\lambda'_{j,t} = \chi h_{j,t}^{\varphi}\left(c_{j,t}-\chi\frac{h_{j,t}^{1+\varphi}}{1+\varphi}\right)^{-\sigma'}$$
$$\lambda'_{j,t} = \left(c_{j,t}-\chi\frac{h_{j,t}^{1+\varphi}}{1+\varphi}\right)^{-\sigma'}$$
$$\exp\left(\varsigma_g\hat{g}_t\right)\lambda'_{j,t} = i_t\beta\mathbb{E}^*_{j,t}\frac{\lambda^c_{j,t+1}}{\pi_{t+1}}.$$

Linearizing each first-order conditions yields:

$$\hat{w}_t = \varphi\hat{h}_{j,t}, \tag{30}$$

[24] Note that our assumptions are not crucial to the model because there exist mechanisms to impose a homogeneous post-saving stock across households. For instance, Andrade et al. (2019) introduce an intra-household risk-sharing plan under which, in each period, there is an agreement between household members that the aggregate amount of bond holdings will be equally shared among them. This is achieved by a transfer plan $z_{j,t}$ within the household to each member $j$ in each period $t$ that are equal to $z_{j,t} = b_{j,t} - B_t/N$. In equilibrium, the sum of transfers is zero $\sum_{j=1}^{N} z_{j,t} = 0$. Hence, in each period, each household has the same post-transfer wealth, even though their consumption levels may differ.

and:

$$\hat{\lambda}'_{j,t} = -\sigma' \left( \bar{c}_j - \chi \frac{\bar{h}_j^{1+\varphi}}{1+\varphi} \right)^{-1} \left( \bar{c}_j \hat{c}_{j,t} - \chi \bar{h}_j^{1+\varphi} \hat{h}_{j,t} \right), \tag{31}$$

$$\hat{\lambda}'_{j,t} = \hat{\imath}_t - \varsigma_g \hat{g}_t + \mathbb{E}^*_{j,t} \left\{ \hat{\lambda}'_{j,t+1} - \pi_{t+1} \right\}. \tag{32}$$

Eq. (30) 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 $\exp(a^{\pi}_{j,t})$ and $\exp(a^{y}_{j,t})$ and, are therefore identical. It thus follows that: $\bar{c}_j = \bar{c}$, $\bar{h}_j = \bar{h}$, $\forall j$.

Hence, we have:

$$\bar{c}\hat{c}_{j,t} - \chi \bar{h}^{1+\varphi} \hat{h}_{j,t} = -\frac{\vartheta}{\sigma'} \left( \hat{\imath}_t - \mathbb{E}^*_{j,t} \hat{\pi}_{t+1} \right) + \frac{\vartheta \varsigma_g}{\sigma'} \hat{g}_t + \mathbb{E}^*_{j,t} \left\{ \left( \bar{c}\hat{c}_{j,t+1} - \chi \bar{h}^{1+\varphi} \hat{h}_{j,t+1} \right) \right\}, \tag{33}$$

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

### A.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 profit $\Pi_{j,t}$ (in Eq. (29)). Hence, we may use the same indexes $j$ and discount factors for firms and household members. Final firms produce a homogeneous good by combining all intermediate goods $\{y_j\}$, $j = 1, \ldots, N$.

**Final sector** The final-good producers are retailers. They buy the intermediate goods and package them into the aggregate supply of goods, denoted by $Y_t^D$, which in equilibrium equals the aggregate good demand from households. In a perfectly competitive market, final producers take the price $P$ of the goods as given and maximize profits as follows:

$$P_t Y_t^D - \sum_{j=1}^{N} p_{j,t} y_{j,t}, \tag{34}$$

subject to a supply constraint:

$$Y_t^D = \left( N^{-1/\epsilon} \sum_{j=1}^{N} {y_{j,t}}^{(\epsilon-1)/\epsilon} \right)^{\epsilon/(\epsilon-1)}. \tag{35}$$

This supply constraint implies that the final-good producers have a technology which aggregates non-perfectly substitutable goods. This imperfect substitutability between all varieties $j$ is driven by the 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 producing price of intermediate firms.

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

$$L = P_t Y_t^D - \sum_{j=1}^{N} p_{j,t} y_{j,t} + \varrho_t \left[ N^{-1/\epsilon} \sum_{j=1}^{N} {y_{j,t}}^{(\epsilon-1)/\epsilon} - \left(Y_t^D\right)^{(\epsilon-1)/\epsilon} \right]. \tag{36}$$

The associated first-order conditions are given by:

$$P_t - \varrho_t (\epsilon-1)/\epsilon \left(Y_t^D\right)^{-1/\epsilon} = 0, \tag{37}$$

$$-p_{j,t} + \varrho_t (\epsilon-1)/\epsilon N^{-1/\epsilon} {y_{j,t}}^{(-1)/\epsilon} = 0, \tag{38}$$

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

$$y_{j,t} = \frac{Y_t^D}{N} \left( \frac{p_{j,t}}{P_t} \right)^{-\epsilon}. \tag{39}$$

The aggregate-price index is given by:

$$P_t = \left[ \frac{1}{N} \sum_{j=1}^{N} p_{j,t}^{1-\epsilon} \right]^{1/(1-\epsilon)}. \tag{40}$$

Applying linearization methods to (39) and (40) leads to the following expressions:

$$\hat{y}_{j,t} = \hat{Y}_t^D - \epsilon \left( \hat{p}_{j,t} - \hat{P}_t \right), \tag{41}$$

$$\hat{P}_t = \frac{1}{N} \sum_{j=1}^{N} \hat{p}_{j,t}. \tag{42}$$

Expressing Equation 42 in growth rates provides the expression for the inflation rate:

$$\hat{\pi}_t = \frac{1}{N} \sum_{j=1}^{N} \hat{\pi}_{j,t} \tag{43}$$

**Intermediate sector**

Firms are homogeneous, distributed on an interval $j \in [1, N]$ and have the following linear production technology:

$$y_{j,t} = h_{j,t}, \tag{44}$$

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

Intermediate-good producers solve a two-stage problem. In the first stage, taking 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 (44).

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

$$\max_{\{y_{j,t}h^d_{j,t}\}} mc_{j,t}y_{j,t} - w_t h^d_{j,t} + \lambda_t \left[h^d_{j,t} - y_{j,t}\right], \tag{45}$$

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:

$$mc_{j,t} = mc_t = w_t. \tag{46}$$

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:

$$ac_{j,t} = \frac{\xi'}{2}\left(\frac{p_{j,t}}{p_{j,t-1}} - \bar{\pi}\right)^2 \frac{Y^D_t}{N}, \tag{47}$$

where $\xi' > 0$ is the price stickiness parameter, $\frac{Y^D_t}{N}$ is the average market share, and $\pi^T$ 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: $y_{j,t} = (p_{j,t}/P_t)^{-\epsilon} Y^D_t/N$. The problem faced by firms is then given by:

$$\max_{\{p_{j,t}\}} \mathbb{E}^*_{j,t} \sum_{\tau=0}^{\infty} \Lambda_{j,t,t+\tau} \left(y_{j,t+\tau}\frac{p_{j,t+\tau}}{P_{t+\tau}} - e^{\varsigma_u \hat{u}_{t+\tau}} mc_{t+\tau}y_{j,t+\tau} - ac_{j,t+\tau}\right), \tag{48}$$

where $\Lambda_{j,t,t+\tau} = \beta^\tau \lambda'_{t+\tau}/\lambda'_t$ is the households' discount factor of agent $j$, $p_{j,t}$ is the individual price set by firm $j$, and $P_t$ is the aggregate price which sets in real terms the problem of firms. Variable $\hat{u}_t$ is an exogenous shock that captures exogenous changes in the cost structure of firms. Parameter $\varsigma_u$ allows to normalize the shock to one in the linearized form of the aggregate-supply equation. Note that the tax rate on the added value, $\tau_{j,t}$, is typically used in the NK literature to offset some market distortions and simplify the analysis of optimal policy. Given the presence of heterogeneity with respect to the benchmark textbook model, the tax rate set by the government is set to offset the relative dispersion in prices $\tau_{j,t} = \frac{p_{j,t}-P_t}{p_{j,t}}$.

Replacing the demand function $y_{j,t} = (p_{j,t}/P_t)^{-\epsilon} Y_t/N$, the objective function of the firms reads as follows:

$$\max_{\{p_{j,t}\}} \mathbb{E}^*_{j,t} \sum_{\tau=0}^{\infty} \Lambda_{j,t,t+\tau} \left( (1-\tau_{j,t+\tau}) \left(\frac{p_{j,t+\tau}}{P_{t+\tau}}\right)^{1-\epsilon} \frac{Y^D_{t+\tau}}{N} - e^{\varsigma_u \hat{u}_{t+\tau}} mc_{t+\tau} \left(\frac{p_{j,t+\tau}}{P_{t+\tau}}\right)^{-\epsilon} \frac{Y^D_{t+\tau}}{N} - ac_{j,t+\tau} \right). \tag{49}$$

The first-order condition reads as:

$$(1-\tau_{j,t}) \frac{(1-\epsilon)}{p_{j,t}} \left(\frac{p_{j,t}}{P_t}\right)^{1-\epsilon} \frac{Y^D_t}{N} + e^{\varsigma_u \hat{u}_t} \epsilon \frac{mc_t}{p_{j,t}} \left(\frac{p_{j,t}}{P_t}\right)^{-\epsilon} \frac{Y^D_t}{N} - \frac{\xi'}{p_{j,t-1}} \left(\frac{p_{j,t}}{p_{j,t-1}} - \bar{\pi}\right) \frac{Y^D_t}{N}$$
$$+ \mathbb{E}^*_{j,t} \Lambda_{j,t,t+1} \frac{p_{j,t+1}}{p^2_{jt}} \xi' \left(\frac{p_{j,t+1}}{p_{j,t}} - \bar{\pi}\right) \frac{Y^D_{t+1}}{N} = 0. \tag{50}$$

Rewriting in terms of firm-specific inflation rates gives:

$$(1-\tau_{j,t})(1-\epsilon) \left(\frac{p_{j,t}}{P_t}\right)^{1-\epsilon} + e^{\varsigma_u \hat{u}_t} \epsilon mc_t \left(\frac{p_{j,t}}{P_t}\right)^{-\epsilon} = \xi' \pi_{j,t} (\pi_{j,t} - \bar{\pi})$$
$$- \mathbb{E}^*_{j,t} \Lambda_{j,t,t+1} \pi_{j,t+1} \xi' (\pi_{j,t+1} - \bar{\pi}) \frac{Y^D_{t+1}}{Y^D_t}. \tag{51}$$

Replacing the government tax on added value yields:

$$\frac{(1-\epsilon)}{\epsilon} \left(\frac{p_{j,t}}{P_t}\right)^{-\epsilon} + e^{\varsigma_u \hat{u}_t} mc_t \left(\frac{p_{j,t}}{P_t}\right)^{-\epsilon} = \frac{\xi'}{\epsilon} \pi_{j,t} (\pi_{j,t} - \bar{\pi})$$
$$- \mathbb{E}^*_{jt} \Lambda_{j,t,t+1} \pi_{j,t+1} \frac{\xi'}{\epsilon} (\pi_{j,t+1} - \bar{\pi}) \frac{Y^D_{t+1}}{Y^D_t}. \tag{52}$$

Applying a linearization of this expression yields to:

$$(\epsilon - 1) \hat{p}_{j,t} + (\epsilon - 1)/\epsilon \left(\varsigma_u \hat{u}_t + \widehat{mc}_t - \epsilon \hat{p}_{j,t}\right) = \frac{\xi'}{\epsilon} \bar{\pi} \hat{\pi}_{j,t} - \frac{\xi'}{\epsilon} \bar{\pi} \beta \mathbb{E}^*_{j,t} \hat{\pi}_{j,t+1}, \tag{53}$$

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

The final inflation equation for each producer $j$ reads as:

$$\hat{\pi}_{j,t} = (\epsilon - 1) \frac{\varphi}{\bar{\pi}\xi'} \hat{h}_t + \beta \mathbb{E}^*_{j,t} \hat{\pi}_{j,t+1} + \hat{u}_t. \tag{54}$$

Note that $\varsigma_u = \xi' \bar{\pi} / (\epsilon - 1)$ normalizes the shock in the linear equation. because the marginal cost is the same across firms, $\widehat{mc}_t = \hat{w}_t$, and from the wage setting, the same across households $\hat{w}_t = \varphi \hat{h}_{j,t}$.

### A.1.4 Authorities

**Monetary policy**. The monetary-policy authority, namely the CB, measures aggregate expectations in the economy by collecting agent forecasts about gross inflation and output, before setting its interest rate. In practice, central banks conduct surveys among firms and households to measure the gap between the relevant target and agents' forecasts. The relevant information here concerns inflation and output, and is measured by an arithmetic average of all of forecasts of inflation and production of the $N$ agents populating the economy:

$$\mathbb{E}_t^* \pi_{t+1} = \frac{1}{N} \sum_{j=1}^{N} \mathbb{E}_{j,t}^* \pi_{j,t+1} \text{ and } \mathbb{E}_t^* Y_{t+1} = \frac{1}{N} \sum_{j=1}^{N} \mathbb{E}_{j,t}^* y_{j,t+1}. \tag{55}$$

The CB sets its interest rate according to a forward-looking Taylor rule subject to an ELB:

$$i_t = \max\left( \bar{\imath} \left( \frac{\mathbb{E}_t^* \pi_{t+1}}{\bar{\pi}} \right)^{\phi_\pi} \left( \frac{\mathbb{E}_t^* Y_{t+1}}{\bar{Y}} \right)^{\phi_y}, 1 \right), \tag{56}$$

where parameters $\bar{\imath}$ , $\bar{\pi}$ and $\bar{Y}$ are long-term values for gross interest, gross inflation rate and average production, respectively. The CB reacts to the deviation of the gross inflation from its steady-state value with a proportion $\phi_\pi$ and $\phi_y$ for output.

The linearized version of the monetary policy rule reads as:

$$\hat{\imath}_t = \max\left( -\bar{r}, \phi_\pi \mathbb{E}_t^* \hat{\pi}_{t+1} + \phi_y \mathbb{E}_t^* \hat{y}_{t+1} \right), \tag{57}$$

with $\bar{r} = \log(\bar{\imath}) = \log(\bar{\pi}/\beta) = \pi^T - \log(\beta)$ and $\hat{\pi}_{t+1}$ the expected inflation rate in the economy between periods $t$ and $t+1$. Note also that linearized average inflation/production forecasts read as:

$$\mathbb{E}_t^* \hat{\pi}_{t+1} = \frac{1}{N} \sum_{j=1}^{N} \mathbb{E}_{j,t}^* \hat{\pi}_{j,t+1} \text{ and } \mathbb{E}_t^* \hat{y}_{t+1} = \frac{1}{N} \sum_{j=1}^{N} \mathbb{E}_{j,t}^* \hat{y}_{j,t+1}, \tag{58}$$

where $E_t^*(\hat{\pi}_{t+1})$ and $E_t^*(\hat{y}_{t+1})$ are the aggregate expectations of the inflation gap and the output gap, respectively.

**Government**. The government implements a value-added tax on firms and borrow $B_t$ from households, while the expenditure side includes interest payments and lump-sum transfers (where all household members receive the same amount). The balance sheet of the government reads as follows:

$$\sum_{j=1}^{N} \tau_{j,t} y_{j,t} + B_t = \sum_{j=1}^{N} T_{j,t} + B_{t-1} i_{t-1}/\pi_t. \tag{59}$$

### A.1.5 General Equilibrium

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

$$\sum_{j=1}^{N} y_{j,t} = \frac{Y_t^D}{N} \sum_{j=1}^{N} \left(\frac{p_{j,t}}{P_t}\right)^{-\epsilon}, \tag{60}$$

which, in a linearized version, reads simply as

$$N\bar{y} \sum_{j=1}^{N} \hat{y}_{j,t} = \bar{Y}^D \left[ N\hat{y}_t^D - \epsilon \left( \sum_{j=1}^{N} \hat{p}_{j,t} - N\hat{P}_t \right) \right]. \tag{61}$$

Using the definition in Equation 42 rules out the effect of price dispersion and allows one to rewrite the general equilibrium as:

$$\frac{1}{N} \sum_{j=1}^{N} \hat{y}_{jt} = \hat{y}_t^D, \tag{62}$$

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

**Final goods sector.** The resource constraint is given by:

$$Y_t^D = \sum_{j=1}^{N} \left( c_{j,t} + \frac{\xi'}{2} (\pi_{j,t} - \bar{\pi})^2 \frac{Y_t^D}{N} \right), \tag{63}$$

where the second term of the right-hand side of Equation 63 is the menu cost stemming from the Rotemberg's price adjustment.

Linearizing Equation 63 yields:

$$\hat{y}_t^D = \frac{1}{N} \sum_{j=1}^{N} \hat{c}_{j,t}, \tag{64}$$

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

**Labor market.** The labor-market equilibrium is reached when aggregate demand from firms satisfies:

$$\sum_{j=1}^{N} h_{j,t}^{d} = \sum_{j=1}^{N} h_{j,t}, \tag{65}$$

which simply reads in a linearized version as :

$$\sum_{j=1}^{N} \hat{h}_{j,t}^{d} = \sum_{j=1}^{N} \hat{h}_{j,t}, \tag{66}$$

since firms and households share the same steady-state values for the number of hours worked.

### A.1.6 Aggregation

**Labor demand dispersion.** In presence of nonseparable preferences, the labor supply in Equation 30 is the same across households:

$$\hat{h}_{j,t} = \hat{h}_t. \tag{67}$$

Combining Equation 44 with the firm-specific demand for intermediate inputs in Equation 41 leads to the following expression

$$\hat{h}_{j,t}^{d} = \hat{y}_{j,t} = \hat{Y}_t^D - \epsilon \left( \hat{p}_{j,t} - \hat{P}_t \right). \tag{68}$$

One can note that if a firm has a pricing strategy that is different from the average – i.e., $\hat{p}_{j,t} \neq \hat{P}_t$ – its demand differs from the total demand – $\hat{y}_{j,t} \neq \hat{Y}_t^D$– and generates a dispersion in demand for labor, as well as output dispersion, across firms. This dispersion in labor demand conflicts with households' homogeneous labor supply in Equation 67.

**Assumption 1** *To map heterogeneous labor demand with homogeneous supply, it is assumed that households evenly split their working time across all firms at no cost, which translates formally as:*

$$h_{j,t} = \sum_{j=1}^{N} \frac{h_{j,t}^{d}}{N} \tag{69}$$

Linearizing this expression from Assumption1 and injecting Equation 68 allows us to express any agent-specific change in labor supply in terms of change in aggregate demand:

$$\begin{aligned}
\hat{h}_{j,t} &= \sum_{j=1}^{N} \frac{\hat{h}^{d}_{j,t}}{N} \\
&= \sum_{j=1}^{N} \frac{\hat{Y}^{D}_{t} - \epsilon\left(\hat{p}_{j,t} - \hat{P}_{t}\right)}{N} \\
&= \hat{Y}^{D}_{t} + \epsilon P_{t} - \epsilon \sum_{j=1}^{N} \frac{\left(\hat{p}_{j,t}\right)}{N} \\
&= \hat{Y}^{D}_{t}
\end{aligned} \tag{70}$$

Even in the presence of heterogeneous demand for labor induced by price dispersion, the general equilibrium assumption in Equation 66 holds. By assuming that households evenly supply their labor to all firms, we rule out the dispersion term and greatly soften the aggregation burden.

**Aggregate demand.** Consider the agent-level linearized Euler equation, Equation 33. Aggregating across all household members leads to:

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

which becomes:

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

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

General equilibrium in labor market allows one to write $\hat{h}_{t} = \hat{h}^{d}_{t} = \hat{y}_{t}$, while general equilibrium in intermediate goods entails $\hat{y}_{t} = \hat{y}^{D}_{t} = \hat{c}_{t}$. 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:

$$\left(1-\chi\right)\hat{y}_{t} = -\frac{\vartheta}{\sigma'}\hat{\imath}_{t} + \mathbb{E}^{*}_{t}\left\{\left(1-\chi\right)\hat{y}_{t+1} + \frac{\vartheta}{\sigma'}\hat{\pi}_{t+1}\right\} + \left(1-\chi\right)\hat{g}_{t}. \tag{71}$$

which is of the same form as the aggregate demand (1) in Section 2.1.

**Aggregate supply.** From the micro-level NK Phillips curve in Equation 54, we may aggregate over all agent types $j$:

$$\sum_{j=1}^{N} \hat{\pi}_{j,t} = \sum_{j=1}^{N} \left[ (\epsilon - 1) \frac{\varphi}{\bar{\pi}\xi'} \hat{h}_t + \beta \mathbb{E}^*_{j,t} \hat{\pi}_{j,t+1} + \hat{u}_t \right],$$

which becomes:

$$\hat{\pi}_t = (\epsilon - 1) \frac{\varphi}{\bar{\pi}\xi'} \hat{h}_t + \beta \mathbb{E}^*_t \hat{\pi}_{t+1} + \hat{u}_t. \tag{72}$$

Note that through a general equilibrium effect, we can rewrite this expression by replacing aggregate labor with aggregate output (as in the usual formulation of the macro textbook) as follows:

$$\hat{\pi}_t = (\epsilon - 1) \frac{\varphi}{\bar{\pi}\xi'} \hat{y}_t + \beta \mathbb{E}^*_t \hat{\pi}_{t+1} + \hat{u}_t. \tag{73}$$

The monetary policy rule is left unchanged, as Eq. (57), which is the same as Eq. (3) in Section 2.1.

### A.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. This step is deemed necessary to make our three-equation NK model similar to the ones found in macro textbooks.

**Marginal utility of consumption.** Let $\lambda_t$ and $\lambda'_t$ denote the marginal utility of consumption under separable and non-separable preferences, respectively. These are given by:

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

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

$$\sigma' = \frac{1-\chi/(1+\varphi)}{1-\chi} \sigma. \tag{74}$$

We may substitute $\sigma'$ by $\sigma$ into Equation 71 as follows:

$$\hat{y}_t = \mathbb{E}^*_t \{\hat{y}_{t+1}\} - \frac{1}{\sigma} \mathbb{E}^*_t \{\hat{\imath}_t - \hat{\pi}_{t+1}\} + \hat{g}_t. \tag{75}$$

**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 New Keynesian Phillips curve:

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

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

$$\xi' = \frac{\varphi}{(\sigma + \varphi)}\xi. \tag{76}$$

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

$$\hat{\pi}_t = (\epsilon - 1)\,\frac{(\sigma + \varphi)}{\bar{\pi}\xi}\hat{y}_t + \beta\mathbb{E}_t^*\hat{\pi}_{t+1} + \hat{u}_t \tag{77}$$

## A.2 Solution under rational expectations

We solve the model under RE using the method of undetermined coefficients (with and without the ELB).

Let us first define the gross and net inflation/interest rates as follows:

$$\pi^T = \log(\bar{\pi}) \text{ and } \bar{r} = \log(\bar{i}) \tag{78}$$

Then, inserting Eq. (3) into Eq. (1) provides the reduced-form expression of the log-linearized model:

$$z_t = \alpha(s_t) + B(s_t)E_t z_{t+1} + \chi^g \widehat{g}_t + \chi^u \widehat{u}_t, \tag{79}$$

with the two endogenous variables $z_t = (\widehat{y}_t\ \widehat{\pi}_t)'$; matrices $\chi^g = (1\ \kappa)'$ and $\chi^u = (0\ 1)'$ are related to the shocks $g$ and $u$ while $\alpha(s_t)$ and $B(s_t)$ are related to the steady-state values and the forward-looking variables, respectively. The values of $\alpha$ and $B$ depend on $s_t$, the state of the CB's policy. The two policy regimes are $s_t = T$ and $s_t = elb$, such that

$$\begin{cases} \widehat{\imath}\left(s_t = T\right)_t = \phi^\pi \mathbb{E}_{j,t}\widehat{\pi}_{t+1} + \phi^y \mathbb{E}_{j,t}\widehat{y}_{t+1} \\ \widehat{\imath}\left(s_t = elb\right)_t = -\overline{r}, \end{cases} \tag{80}$$

The transition between the two regimes is given by matrix $Q$

$$Q = \begin{bmatrix} q^T & 1 - q^T \\ 1 - q^{elb} & q^{elb} \end{bmatrix} \tag{81}$$

Given Eq. (79), the general form of the MSV solution reads as:

$$z_t = a(s_t) + c(s_t)\widehat{g}_t + d(s_t)\widehat{u}_t, \tag{82}$$

where the coefficient values in matrices $a$, $c$ and $d$ depend on whether the ELB is binding or not.

Taking expectations based on an AR(1) specification of the stochastic processes $\widehat{g}$ and $\widehat{u}$ with autoregressive coefficients $\rho^g$ and $\rho^u \in (0,1)$ yields (assuming for now that shocks are observable in $t$):

$$\begin{cases} E\left(z\left(s_t=T\right)_{t+1}\right)= & q^T\{\alpha\left(s_t=T\right)+g_t\left(c\left(s_t=T\right)\rho^g\right)+u_t\left(d\left(s_t=T\right)\rho^u\right)\}... \\ & +\left(1-q^T\right)\{\alpha\left(s_t=elb\right)+g_t\left(c\left(s_t=elb\right)\rho^g\right)+u_t\left(d\left(s_t=elb\right)\rho^u\right)\} \\ E\left(z\left(s_t=elb\right)_{t+1}\right)= & \left(1-q^{elb}\right)\{\alpha\left(s_t=T\right)+g_t\left(c\left(s_t=T\right)\rho^g\right)+u_t\left(d\left(s_t=T\right)\rho^u\right)\}... \\ & +q^{elb}\{\alpha\left(s_t=elb\right)+g_t\left(c\left(s_t=elb\right)\rho^g\right)+u_t\left(d\left(s_t=elb\right)\rho^u\right)\}. \end{cases} \tag{83}$$

Given that shocks are *i.i.d.* $\rho^g=\rho^u=0$, there is no source of persistence in the model and expectations in their rational form are not history dependent. Thus the MSV reduces to a vector of intercepts $E\left(z_{t+1}\right)=a(s_t)$. Given the fact that the Taylor rule is driven by $t+1$ expectations, the interest rate is static. Thus, there is no possibility for the interest rate to move from one regime to the other and the ELB never binds under RE. Accordingly, there is no probability of switching between the two regimes under RE, which is equivalent to

$$Q=\begin{bmatrix} q^T=1 & 0 \\ 0 & q^{elb}=1 \end{bmatrix}. \tag{84}$$

Expectations boil down to:

$$\begin{cases} E\left(z\left(s_t=T\right)_{t+1}\right)= & a^T, \\ E\left(z\left(s_t=elb\right)_{t+1}\right)= & a^{elb}. \end{cases} \tag{85}$$

It is possible to verify the values of $Q$ by plugging the model into an occasionally-binding-constraint-solution algorithm that endogenises the value of $Q$. We do so by using Occbin from Guerrieri & Iacoviello (2015).[25]

Inserting Eq. (85) back into (79) uniquely identifies the MSV solution as:

$$z_t=\alpha\left(s_t\right)+B\left(s_t\right)a\left(s_t\right)+g_t\chi^g+u_t\chi^u, \tag{86}$$

with $a(s_t)=(I-B(s_t))^{-1}\alpha(s_t)$, which makes clear that the coefficient values of matrices $B(s_t)$ and $\alpha(s_t)$ depend on the regime $s_t$

In this specific case, RE are either anchored to the targeted steady state or to the unintended steady-state. Under the occasional-binding-constraint logic, the targeted steady state is the initial state in the model. We consider the REE at the targeted steady state

[25]Simulations are available upon request.

that we denote by a 'T' superscript (for 'target'). We insert the specification of the Taylor rule (3) when the ELB is not binding ($\widehat{\imath}_t = \phi^\pi \mathbb{E}_t \widehat{\pi}_{t+1} + \phi^y \mathbb{E}_t \widehat{y}_{t+1}$) into (1) and obtain the expressions $\alpha^T = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and $B^T = \begin{bmatrix} 1-\sigma^{-1}\phi^y & \sigma^{-1}(1-\phi^\pi) \\ \kappa(1-\sigma^{-1}\phi^y) & \beta + \sigma^{-1}(1-\phi^\pi)\kappa \end{bmatrix}$.

The MSV-REE solution at the target is then given by:

$$a^T = (I - B^T)^{-1} \alpha^T \tag{87}$$

Similarly, when the ELB is binding, the monetary policy rule reads as $\widehat{\imath}_t = -\bar{r}$. Inserting this expression back into Eq. 1, the REE at the ELB, which we denote with an *elb* superscript, is described by:

$$\alpha^{elb} = \begin{bmatrix} \sigma^{-1}\bar{r}, \ \kappa\sigma^{-1}\bar{r} \end{bmatrix} \text{ and } B^{elb} = \begin{bmatrix} 1 & \sigma^{-1} \\ \kappa & \beta + \sigma^{-1}\kappa \end{bmatrix}. \tag{88}$$

The MSV-REE solution at the ELB is then given by:

$$a^{elb} = (I - B^{elb})^{-1} \alpha^{elb} \tag{89}$$

To simulate the model under RE, we assume that expectations are anchored to the targeted steady state so that the ELB is never binding.

## A.3 Determinacy and E-stability

The REE (89) is determinate under RE if the two eigenvalues of matrix $B^T$ lie within the unit circle. This is the case if all three conditions $\phi^y < \sigma(1+\beta^{-1})$, $0 < \kappa(\sigma_\pi - 1) + (1+\beta)\sigma_y < 2\sigma(1+\beta)$ and $\kappa(\phi^\pi - 1) + (1-\beta)\phi^y > 0$ hold (Bullard & Mitra 2002, p. 1121). Our calibration imposes these restrictions on the parameters' values. Specifically, the REE values at the target are $a^T = (00)'$, and the REE is determinate (the two eigenvalues are complex and equal $\lambda_i^- = .933 - .027i$ and $\lambda_i^+ = .933 + .027i$). Note that the same conditions ensure that this solution is E-stable, *i.e.,* stable if agents use adaptive learning instead of RE (Bullard & Mitra 2002).

By contrast, the REE at the ELB (88) is indeterminate under RE and unstable under learning. To see this, notice that the characteristic polynomial of $B^{elb}$ is $\beta + \lambda(-1-\beta-\kappa\sigma^{-1}) + \lambda^2 = 0 \Leftrightarrow a_0 + a_1\lambda + \lambda^2$. For both eigenvalues to be within the unit circle and the REE to be determinate, we need $\mid a_0 \mid < 1$ and $\mid a_1 \mid < 1 + a_0$. The first condition always holds as $\beta < 1$, but the second is always violated as $\sigma^{-1}\kappa > 0$. Therefore, the deflationary state is indeterminate under RE and features multiple equilibria.[26]

[26] Another way to see this is to note that the ELB corresponds to an interest rate peg that is known to give rise to indeterminacy.

Furthermore, the determinant of $B^{elb} - I$ ($I$ being the identity matrix) is $-\sigma^{-1}\kappa < 0$, which implies that one eigenvalue of $B^{elb} - I$ has a negative real part and one has a positive real part (equivalently, one eigenvalue of $B^{elb}$ is lower than one, the other is not). Therefore, under learning, the deflationary steady state is unstable and is a saddle.

Under our calibration, the REE values at the ELB are $a^{elb} = (-0.007 - 0.013)'$, and the two eigenvalues of $B^{elb}$ are real and equal $\lambda_i^- = 0.906 < 1$ and $\lambda_i^+ = 1.099 > 1$.

## A.4 Calculating the welfare cost

In this section, we develop the approximation of the welfare criterion and use it to compute the welfare cost of fluctuations.

## A.5 Welfare function

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

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

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, $h_{j,t} = Y_t/N$. This means we can rewrite $h_{j,t}$ 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 $\gamma_{j,t} = Nc_{j,t}/C_t$, with $C_t = \sum_{j=1}^{N} c_{j,t}$. The resulting utility of agent $j$ is given by:

$$U_{j,t} = \frac{1}{1-\sigma'}\left(\gamma_{j,t}\frac{C_t}{N} - \chi \frac{(Y_t/N)^{1+\varphi}}{1+\varphi}\right)^{1-\sigma\prime}.$$

To take into account the role of nominal ridigities, we use the resource constraint in Equation 63 to replace aggregate output as follows: $C_t = Y_t^D\left(1 - \frac{1}{N}\sum_{j=1}^{N}\frac{\xi'}{2}\left(\pi_{jt} - \bar{\pi}\right)^2\right)$. In addition, we express output in supply terms using the general-equilibrium condition in the intermediate sector in Equation 60: $Y_t^D = Y_t\left(\frac{1}{N}\sum_{j=1}^{N}\left(\frac{p_{j,t}}{P_t}\right)^{-\epsilon}\right)^{-1}$. 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:

$$U_{j,t} = \frac{1}{1-\sigma'}\left(\gamma_{j,t} y_t \frac{\left(1-\frac{1}{N}\sum_{j=1}^{N}\frac{\xi'}{2}\left(\pi_{j,t}-\bar{\pi}\right)^2\right)}{\left(\frac{1}{N}\sum_{j=1}^{N}\rho_{j,t}^{-\epsilon}\right)} - \chi\frac{y_t^{1+\varphi}}{1+\varphi}\right)^{1-\sigma\prime},$$

where $\rho_{j,t} = p_{j,t}/P_t$ and $y_t = Y_t/N$.

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:

$$E_t\left(U_{j,t}\right) \simeq \bar{U}_j + \frac{U_{\gamma\gamma}}{2} var_t\left(\gamma_{j,t}\right) + \frac{U_{yy}}{2} var_t\left(y_t\right) + N\frac{U_{\rho\rho}}{2} E_j var_t\left(\rho_{j,t}\right) + N\frac{U_{\pi\pi}}{2} E_j var_t\left(\pi_{j,t}\right),$$

where derivatives are computed from a symbolic toolbox.

Note that $E_j var_t\left(\pi_{j,t}\right)$ is the average volatility of inflation across firms, while $E_j var_t\left(\rho_{j,t}\right)$ 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:

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

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

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

$$E_t E_j\left(U_{j,t}\right) \simeq \frac{1}{N}\sum_{j=1}^{N}\bar{U}_{j,t}, \tag{90}$$

which can be expressed as:

$$E_t E_j\left(U_{j,t}\right) \simeq \bar{U}_j + \frac{U_{\gamma\gamma}}{2} E_j var_t\left(\hat{\gamma}_{j,t}\right) + \frac{U_{yy}}{2} var_t\left(\hat{y}_t\right) \tag{91}$$

$$+ N\frac{U_{\rho\rho}}{2} E_j var_t\left(\hat{\rho}_{j,t}\right) + N\frac{U_{\pi\pi}}{2\bar{\pi}^2} E_j var_t\left(\hat{\pi}_{j,t}\right). \tag{92}$$

Average utility is connected to average welfare as follows:

$$E_t \mathcal{W}_t = \frac{E_t E_j \left(U_{j,t}\right)}{1-\beta}, \tag{93}$$

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

$$E_t\left(U_t\right) \simeq \bar{U} + \frac{U_{yy}}{2} var_t\left(\hat{y}_t\right) + N\frac{U_{\pi\pi}}{2\bar{\pi}^2} var_t\left(\hat{\pi}_t\right). \tag{94}$$

# B Moment matching

We proceed by building on the related literature on the estimation of macroeconomic models using simulated moment matching methods; see (Hansen 1982, McFadden 1989). This method aims to infer the values of structural parameters by minimizing the squared distance between the moments simulated by a model and their observable counterparts.

## B.1 The SMM estimator

Under the SL process, the model – described by Eq. 1 to 3 combined with the SL algorithm introduced in Section 2.3 – can be expressed in the following compact form:

$$\mathbb{E}_t^{SL}\left\{f_\Theta\left(z_{t+1}, z_t, \varepsilon_t\right)\right\} = 0, \tag{95}$$

where $z_t$ is the set of endogenous variables, $\varepsilon_t$ the set of *i.i.d.* shocks and $f_\Theta(\cdot)$ the equations of the model based on the set of structural parameters $\Theta$. First, we partition the parameters $\Theta$ into two subsets. The first set contains the calibrated parameters as given in Table 1. The second set, $\theta \in \Theta$, contains the parameters that are determined by SMM method.

The SMM estimator is defined as:

$$\hat{\theta}_{SMM} = \arg\min_\theta \left[m_T\left(x_t\right) - m_{s,\tau}\left(x_t^\theta\right)\right]' W \left[m_T\left(x_t\right) - m_{s,\tau}\left(x_t^\theta\right)\right], \tag{96}$$

where $m_T\left(x_t\right) - m_{s,\tau}\left(\hat{x}_t^\theta\right)$ is the distance vector between the observed and the simulated moments that we seek to minimize, and $W$ is a weighting matrix. Hence, the matrix product in Eq. (96) provides the sum of the squares of the residuals between the observed and matched moments. Let us now define $m_T\left(x_t\right)$, a $p \times 1$ vector of moments calculated using stationary and ergodic real data $x_t$ of sample size $T$, and $m_{s,\tau}\left(x_t^\theta\right)$. The model-generated counterpart based on artificial series $x_t^\theta$ using the set of parameters $\theta$ fed in (95), while $s$ is the number of parallel chains (or Monte-Carlo iterations) that we draw to compute the moments with a sample size $\tau$.

With respect to our optimization problem, we set $p = 10$ as we match exactly 10 moments. Since our sample has a size $T = 200$ (quarters), we impose a similar size for our simulated data $T = \tau = 200$. We also fix the number of parallel chains to $s = 100$. As some specific sets of parameters combined with specific realizations of shocks may create explosive dynamics for some parallel draws, we discard these chains when computing the moments, and ensure that this number of chain never exceeds 5% of the simulations. Finally, regarding the weighting matrix $W$, we impose the identity matrix as $\hat{\theta}_{SMM}$ is consistent with any positive-definitive weighting matrix (see Ruge-Murcia (2012) for a discussion on this aspect).

Note that the mutation process of the SL algorithm represents an additional source of stochasticity beyond the fundamentals shocks (in this case, cost-push and real shocks).

To deal with this new source of noise, we follow Arifovic et al. (2013, 2018) and take the median simulation across the $n$ different series of mutations. Doing so means that the computation of the objective function in (96) would require computing $n * s * \tau$ values of output and inflation.[27] We therefore face a heavy computational burden for the inference of the structural parameters, which we alleviate using two different assumptions: controlling the noise arising from the SL process and adding information on priors. We discuss these two aspects in the following subsections.

## B.2 Managing the noise from the mutations

To reduce the computational burden, we first control the noise in mutations by selecting the mutation sequence that generates the simulations that is the closest to the median. Specifically, we draw $s = 100$ chains of shocks $u$ and $g$ at the beginning of the matching procedure and keep them unchanged. For each of the $s$ chains, we run $n = 100$ Monte Carlo simulations of the model under SL and only retain one representative simulation. To select this representative simulation, we choose the one for which the squared distances of inflation and output gaps to their median values over the 100 replications is the smallest.[28] Therefore, for each chain of shocks, we retain only one simulation. We do so for each of the $s = 100$ series of shocks. This shrinks the computational burden by a factor $n = 100$ without affecting the contribution of SL to the dynamics of output and inflation.

## B.3 Adding information on priors

Because of the computational burden, we propose to include information about priors in order to reduce computing times, in the same spirit as Ruge-Murcia (2012). Ruge-Murcia (2012) proposes a mixed estimation approach characterized by a prior information that aims to avoid the exploration of some parameter spaces that are economically irrelevant. We treat priors as additional moments to match in the objective function. We denote by using $\mathcal{P}(\theta)$ the sum of the pdf stemming from the priors for $\theta$. The resulting quasi-Bayesian SMM estimator is then defined as:

$$\hat{\theta}_{SMM} = \arg\min_{\theta} \left[m_T(x_t) - m_{s,\tau}\left(\hat{x}_t^{\theta}\right)\right]' W \left[m_T(x_t) - m_{s,\tau}\left(\hat{x}_t^{\theta}\right)\right] + \Xi\mathcal{P}(\theta), \tag{97}$$

[27] For example, imposing $n = 100$ as in Arifovic et al. (2018) would necessitate computing 100*100*200=2e6 artificial values of output and inflation. For that parametrization, a rational expectations model only requires 100*200=2e4 values to compute the moments.

[28] Note that in a one-dimensional problem, this procedure boils down to selecting the median. However, as we match a two-dimensional model (inflation and output gaps), our procedure provides a way to approximate the median simulation.

The first term of the expression is the same as in (96), the second term $\Xi\mathcal{P}(\theta)$ introduces a penalty into the objective function when the matched values range differ from their prior distributions. Here, $\Xi$ is the relative weight of prior information with respect to the squared distance of moments. For loglikelihood-based methods, $\Xi = 1$ when the number of observations is high. In contrast, here we only have 10 observations, which makes our prior information dominate the objective function. We set the weight on priors to $\Xi = 1/s = 0.01$ in order to mimic the decreasing relative weight of priors when using full information methods.[29]

## B.4 Optimization

We solve Eq. (96) using the *CMAES* optimization algorithm of Hansen et al. (2003). The CMAES algorithm is a global optimization strategy that has the ability to deal with large-scale optimization problems and avoid local minima. This algorithm provides an accurate measure of the Hessian matrix, even in the presence of bound restrictions and priors for control variables, as is the case in Eq. (96). Specifically, learning the covariance matrix in the CMAES is analogous to learning the inverse Hessian matrix in a gradient-based, local optimization method such as the quasi-Newton method.

[29]In the full information estimation case, the weight of prior is decreasing in the number of observations.

# C An alternative heterogeneous-expectations model

This section presents an alternative heterogeneous-expectation model and its calibration based on a moment-matching method.

## C.1 Heterogeneous expectations under adaptive learning

This model is a version of the heterogeneous-expectation NK model of Branch & McGough (2010) and Hommes & Lustenhouwer (2019*b*). Aggregate expectations, denoted by a star-superscript, are the weighted sum of two forecasting strategies, labeled simply by 1 and 2, such that:

$$E_t^* \left\{ \hat{z}_{t+1} \right\} = \begin{bmatrix} E_t^* \left\{ \hat{y}_{t+1} \right\} \\ E_t^* \left\{ \hat{\pi}_{t+1} \right\} \end{bmatrix} = \begin{bmatrix} n_t^y a_{1,t-1}^y + (1 - n_t^y) a_{2,t-1}^y \\ n_t^\pi a_{1,t-1}^\pi + (1 - n_t^\pi) a_{2,t-1}^\pi \end{bmatrix}, \tag{98}$$

with $n_t^x$ and $1 - n_t^x$, $x = \{\pi, y\}$, the respective shares of agents using strategies 1 and 2 for forecasting the inflation and output gaps. Under i.i.d disturbances, shocks are irrelevant for forecasting $t+1$ variables, and agents only use the intercept to form their forecasts. Thus, in line with our implementation of the model under SL, the learning process corresponds to 'steady-state learning' (Evans et al. 2008), and the two strategies, 1 and 2, to forecast variable $x = \{y, \pi\}$ are given by:

$$\begin{bmatrix} a_{1,t}^x \\ a_{2,t}^x \end{bmatrix} = \begin{bmatrix} a_{1,t-1}^x + \eta_1 (a_{1,t-1}^x - \hat{x}_t) \\ a_{2,t-1}^x + \eta_2 (a_{2,t-1}^x - \hat{x}_t) \end{bmatrix}, \tag{99}$$

where the heterogeneity of forecasts stems from differences in the gain parameter associated which each strategy, namely $0 \leq \eta_1^x, \eta_2^x < 1$, with $\eta_1^x \neq \eta_2^x$ for each variable $x$. If $\eta_1^\pi = \eta_2^\pi = \eta_1^y = \eta_2^y$, the model boils down to the standard adaptive learning model with homogeneous expectations.

Based on Brock & Hommes (1997), the relative share of each strategy in the population evolves according to a heuristic-switching model where the best performing strategy tends to be adopted by a larger share of agents at the expense of the worst performing one. There are two heuristic-switching models, one for each forecast $\pi$ and $y$, in line with our implementation of SL, where we use two separate tournaments for the two variables.[30] Abstracting from the time index for a moment, let the share of agents $n^x$ using strategy 1 to forecast variable $x$ be:

$$n^x = \frac{\exp(\omega^x R_1^x)}{\exp(\omega^x R_1^x) + \exp(\omega^x R_2^x)}, \tag{100}$$

[30]This way, the number of free parameters to be matched against the empirical moments will be conveniently the same under this model and the SL model.

with $R_1^x$ and $R_2^x$ the fitness of the forecasting strategies 1 and 2 respectively, and $\omega^x \geq 0$ the so-called 'intensity of choice' related to the choice of a rule to forecast variable $x$. Let us provide some intuition on the functioning of Eq. (100).

The share of strategy 1 in the forecasting rules of variable $x$, $n^x$, is an increasing function of $R_1^x$ (or a decreasing function of $R_2^x$): if strategy 1 has a higher fitness than strategy 2, more agents will switch to use Strategy 1, and vice-versa if strategy 2 is the out-performer. The right-hand-side of Eq. (100) is bounded between 0 and 1 such that $n_z^x$ in line with the interpretation of $n^x$ as a share of the agents. The higher the intensity of choice is, the more likely agents are to switch to the best-performing strategy. In the limiting case where the intensity of choice equals $\omega^x = 0$, each strategy is chosen with equal probability, independently of its performance. In the opposite case where $\omega^x = \infty$, the best-performing strategy is always chosen by all agents in each period.

We now define $F^x = R_1^x - R_2^x$ as the difference between the fitness of strategies 2 and 1. Following, e.g., Branch & Gasteiger (2019), and dropping the $x$-superscript for readability, one can write:

$$n = \frac{\exp(\omega R_1)}{[\exp(\omega R_1) + \exp(\omega R_2)]} = \frac{\exp(\omega F)}{[\exp(\omega F) + 1]}, \tag{101}$$

and likewise (102)

$$1 - n = \frac{\exp(-\omega F)}{[1 + \exp(-\omega F)]}. \tag{103}$$

Taking the difference between $n$ and $1 - n$, we obtain:

$$n - (1 - n) = 2n - 1 = \frac{\exp(\omega F)}{[\exp(\omega F) + 1]} - \frac{\exp(-\omega F)}{[1 + \exp(-\omega F)]} \tag{104}$$

$$\Leftrightarrow n = (1/2) \times \{\tanh[(\omega/2) \times F] + 1\}. \tag{105}$$

Applying Eq. (104) to our two-dimensional forecasting model and reintroducing the time dimension, we obtain the laws of motion of the fractions of agents using strategy 1 for forecasting variables $y$ and $\pi$ as:

$$\begin{bmatrix} n_t^y \\ n_t^\pi \end{bmatrix} = \begin{bmatrix} (1/2) \times \{\tanh[(\omega^y/2) \times F_{t-1}^y] + 1\} \\ (1/2) \times \left\{\tanh\left[(\omega^\pi/2) \times F_{t-1}^\pi\right] + 1\right\} \end{bmatrix}, \tag{106}$$

and the fractions using strategy 2 are simply given by $1 - n_t^y$ and $1 - n_t^\pi$.

What remains to be defined is a measure of a strategy's fitness. As usual in the related literature, and in line with the implementation under SL, we assume that a strategy fitness is an increasing function of its relative forecast accuracy with respect to the other strategy. Specifically, the fitness of strategy $j = 1, 2$ in forecasting variable $x$ is given by the weighted sum of the squared past forecast errors available in time $t$:

$$\left[\begin{array}{c} R^y_{j,t} \\ R^\pi_{j,t} \end{array}\right] = \left[\begin{array}{c} (1-\delta^y)R^y_{j,t-1} - \delta^y(y_{t-1} - a^y_{j,t-1})^2 \\ (1-\delta^\pi)R^\pi_{j,t-1} - \delta^\pi(\pi_{t-1} - a^\pi_{j,t-1})^2 \end{array}\right], \tag{107}$$

where $1 > \delta^x \geq 0$ is a memory coefficient which introduces history dependence in the forecasting performance measures. If $\delta^x = 0$, $n^x_t$ is only a function of the relative forecast error in period $t-1$. In the opposite case, when $\delta^x \to 1$, $n^x_t$ becomes the sum of all past forecasting errors.

## C.2 Moment matching with the alternative expectations model

We apply the same moment-matching method to determine the values of the learning parameters as in Section 3 with the SL model. As for the SL model, a vector $\Theta = [\eta_1, \eta_2, \omega^y, \omega^\pi, \delta^y, \delta^\pi]$ gathers the six free parameters that define the law of motion of the expectations. The same moments as in the SL model are employed to match the data.

The priors are set such that one of the forecasting strategies has a larger gain than the other one. This asymmetry allows for the possibility of one rule coming to represent unstable trend-chasing behavior, while the other (with a gain off less than one) captures stabilizing mean-reverting behavior. Thus, the learning gain $\eta_1$ has a prior set with an Inverse-Gamma distribution of mean 0.01 and standard-deviation 0.02, while we set the prior of $\eta_2$ using a Weibull distribution of mean 1 and standard deviation 0.5. The difference between the two gains tunes the heterogeneity of expectations and is essential for matching the moments from the cross-section in the SPF forecasts.

The fitness-decay rates use a standard beta prior distributions bounded by values between 0 and 1. The intensity-of-choice prior means are set at 5, based on the estimates of the heuristic-switching model on SPF data in Cornea-Madeira et al. (2019). The distributions are assumed to be normal with a relatively large standard deviation of 10.

It is well-known in the learning literature that learning models exhibit unstable dynamics as long as the ELB binds for too long (Evans et al. 2008). The present model is no exception: with an ELB, this heterogeneous-expectation model diverges along deflationary spirals in 69% of the runs using the same calibrated parameters and the same chains of shocks that we used to calibrate the SL model with an ELB in Section 3. To address this instability issue under adaptive learning, one solution consists of complexifying the model, along the lines of Ozden (2021). They consider a regime-switching implementation, on top of the heuristic-switching model, where it is assumed that some of the agents have expectations that always remain anchored at the target. In a different fashion, Evans et al. (2022) introduce an exogenous lower bound on output to introduce a third stable steady state with deflation. This lower bound is, however, ad-hoc. A third solution is to take a shortcut and implement the model under shadow rates, namely assuming that the ELB is irrelevant, for instance because unconventional monetary policy can achieve economic outcomes comparable to negative nominal rates. However, there is ample evidence that the ELB is relevant for business cycle dynamics (Aruoba et al. 2022). The fourth option is to use a model such as our SL model that does not suffer from this issue: it is simulated with a constrained monetary policy.

None of the solutions named above are perfect. Because our primary objective is to keep the alternative model simple and free from ad-hoc mechanisms in order to compare it with our SL results, we calibrate the heterogeneous-expectations model with adaptive learning under shadow rates. Results of such a moment-matching procedure are reported in Tables 7 and 8.

Even under shadow-rate setting, the AL model fails to produce the crucial heterogeneity in inflation expectations. Admittedly, the AL model without an ELB performs better overall than the SL model with an ELB, but most of the gain comes from the closer match of the autocorrelation of the real variables, which may not realistically be accounted for with white-noise shocks only. As explained in Section 2, we instead see the proportions of the autocorrelation in inflation and output gaps that are matched by a learning model as a proxy for the share of persistence in the economy that may arise due to deviations from RE. It does not seem plausible that *all* persistence, or even most of it, arises from such a source. The optimized HAL model also failed to account for the small but positive correlation between inflation and output.

In conclusion, despite relaxing the ELB constraint, which gives an obvious advantage to the alternative model, our SL model i) achieves a similar fit of the data, ii) does a better job of matching the heterogeneity of inflation forecasts, which is a crucial target for CB communication, while iii) including the ELB constraint, iv) remaining simple and v) without assuming any arbitrary bound on expectations or endogenous variables.

| | | Simulated | Empirical *[0.005;0.995]* |
|---|---|---|---|
| $\sigma(\hat{y}_t)$ | output gap sd. | 4.37 | 4.38 *[3.97;4.83]* |
| $\rho(\hat{y}_t, \hat{y}_{t-1})$ | output gap autocorr. | 0.68 | 0.98 *[0.98;0.99]* |
| $\sigma(\hat{\pi}_t)$ | inflation gap sd. | 0.64 | 0.60 *[0.54;0.66]* |
| $\rho(\hat{\pi}_t, \hat{\pi}_{t-1})$ | inflation gap autocorr. | 0.92 | 0.90 *[0.87;0.92]* |
| $\rho(\hat{\pi}_t, \hat{y}_t)$ | inflation-output correlation | -0.02 | 0.08 *[-0.07;0.21]* |
| $\Delta^{\mathbb{E}y}$ | av. forecast dispersion of output gap | 0.30 | 0.36 *[0.31;0.41]* |
| $\Delta^{\mathbb{E}\pi}$ | av. forecast dispersion of inflation gap | 0.05 | 0.25 *[0.22;0.28]* |
| $\rho(\Delta_t^{\mathbb{E}y}, \Delta_{t-1}^{\mathbb{E}y})$ | autocorr. of forecast disp. of output gap | 0.58 | 0.76 *[0.70;0.82]* |
| $\rho(\Delta_t^{\mathbb{E}\pi}, \Delta_{t-1}^{\mathbb{E}\pi})$ | autocorr. of forecast disp. of inflation gap | 0.56 | 0.64 *[0.55;0.72]* |
| $P[\hat{\imath}_t > -\bar{r}]$ | probability not at the ELB | 0.87 | 0.86 *[0.81;0.91]* |
| Objective function value | | 0.29 | $\times$ |

Table 7: Comparison of the (matched) theoretical moments with their observable counterparts

| | | | Prior Distributions | | | Posterior Results | |
|---|---|---|---|---|---|---|---|
| | | | Shape | Mean | STD | Mean | STD |
| $\sigma_g$ | - | real shock std | Invgamma | .1 | 5 | 2.6552 | 0.4607 |
| $\sigma_u$ | - | cost-push shock std | Invgamma | .1 | 5 | 0.18946 | 0.0200 |
| $100\pi^T$ | - | quarterly inflation trend | Beta | .62 | .1 | 0.8 | 0.4607 |
| $\kappa$ | - | Phillips curve slope | Beta | .05 | .1 | 0.013141 | 0.0041 |
| $\eta_1$ | - | gain for PLM 1 | Invgamma | .01 | .02 | 0.0059005 | 0.0055 |
| $\eta_2$ | - | gain for PLM 2 | Weibull | 1 | .5 | 0.80085 | 0.0737 |
| $\omega^y$ | - | intensity of choice for $\mathbb{E}y$ | Normal | 5 | 10 | 0.66514 | 0.1859 |
| $\omega^\pi$ | - | intensity of choice for $\mathbb{E}\pi$ | Normal | 5 | 10 | 17.139 | 2.2049 |
| $\delta^y$ | - | fitness decay rate for $\mathbb{E}y$ | Beta | .5 | .2 | 0.074572 | 0.0143 |
| $\delta^\pi$ | - | fitness decay rate for $\mathbb{E}\pi$ | Beta | .5 | .2 | 0.47755 | 0.1643 |

Notes: The values of the standard deviation of the optimized parameters are computed using a numerical approximation of a sparse matrix representation of the Hessian matrix.

Table 8: Optimized parameters using the simulated method of moments matching the SPF data (1968–2018)