De-anchored Inflation Expectations and Monetary Policy*

Johannes J. Fischer†

November, 2023

Abstract

This paper studies the conduct of monetary policy in a model with an endogenous degree of expectations anchoring. I use an estimated New Keynesian model with endogenous forecast switching to replicate the time-varying excess sensitivity of long-term inflation expectations to inflation surprises as well as the resulting movements of long-term inflation expectations. In this model, de-anchoring leads to increased inflation volatility and can cause deflationary spirals when the zero lower bound (ZLB) is binding. This implies an optimal inflation response that is significantly higher when expectations are endogenously anchored compared to rational expectations. Price Level Targeting, on the other hand, can increase the risk of deflationary spirals near the ZLB.

JEL Classification: E31, E52, E70

Keywords: Monetary policy, Inflation Expectations, Non-rational Expectations

*I am indebted to my supervisors Ramon Marimon and Edouard Challe for their continuous guidance and support. I, furthermore, thank Leonardo Melosi, Martin Hellwig, Mirko Wiederholt, Fabio Milani, Michael Bauer, Matthias Rottner, Oliver Pfäuti, Wolfram Horn, as well as the participants at various seminars and conferences for helpful comments.

†Bank of England; e-mail: johannes.fischer@bankofengland.co.uk

1. Introduction

The anchoring of inflation expectations is a central tenet of monetary policy making. Most central banks in advanced economies try to anchor private long-term inflation expectations to their inflation target. Anchored long-term expectations reflect the trust in the central bank's commitment to offset inflationary shocks and to bring inflation eventually back to target. However, when households lose trust in the central bank's commitment, they start to believe that temporary inflation movements will not be offset by the central bank. This may cause long-term expectations to eventually de-anchor from the central bank's target.[1] The resulting pass-through of short-term conditions to long-term expectations can lead to self-reinforcing inflation movements as in the 1970s (see Figure 1). [2]. More recently, fears of expectations de-anchoring have resurfaced as consequence of both the low inflation rates of the 2010s as well as the high inflation rates of in the aftermath of the Covid-19 pandemic (Reis, 2021) and the war in Ukraine.

Therefore, this paper studies the conduct of monetary policy when the anchoring of expectations is determined endogenously. I show that de-anchoring can occur if households do not believe in the central bank's ability to stabilise inflation even in the face of transitory shocks and a central bank that is committed to stabilise inflation.[3] I take the model to the data, providing a general equilibrium model with endogenous de-anchoring that is fully estimated using likelihood methods and expectations data. The estimated model implies a substantial degree of expectations de-anchoring beginning in 2021. Using this model, I evaluate two policy frameworks, Inflation Targeting (IT) and Price Level Targeting (PLT), in terms of their stability properties and their ability to stabilise inflation.

This delivers three key policy recommendations: First, expectations de-anchoring at the Zero Lower Bound (ZLB) can lead to deflationary spirals, whereas expectations de-anchoring away from the ZLB does not lead to inflationary spirals. This implies an asymmetric welfare

[1] There is no widely agreed-upon definition of expectations de-anchoring (Kumar et al., 2015). Instead, researchers work with a set of properties commonly thought to characterise de-anchored expectations, such as sensitivity of long-term expectations to inflation news or short-term expectations. A common theme of de-anchoring is that households perceive temporary shocks to have persistent effects, even if this is not justified by the fundamental structure of the economy. Because these expectations are not necessarily model-consistent, de-anchoring is not easily handled in the context of rational expectations, as Ben Bernanke pointed out in his speech at the NBER Monetary Economics Workshop in July 2007.

[2] To put it differently, anchored expectations can reduce the endogenous amplification of shocks. Among other things, this result shows up in the flatter reduced-form New Keynesian Phillips curve of (Hazell et al., 2022)

[3] In this paper I study the possibility of de-anchoring despite the central bank's commitment to stable inflation. Of course, de-anchoring can also arise when the central bank does not attempt to stabilise inflation, e.g. because it misinterprets the state of the economy or because it accommodates the fiscal authority. In this case policy prescriptions are more straightforward compared to the case when de-anchoring is driven by agents' beliefs.

cost of expectations de-anchoring. Second, near the ZLB the optimal inflation response is significantly higher when expectations are endogenously anchored compared to rational expectations. Third, Price Level Targeting near the ZLB can actually de-stabilise the economy unless the central bank reacts extremely aggressive to movements of inflation.

The key feature of my model is that a representative household weighs his beliefs about whether or not the Taylor principle is satisfied: if he trusts the central bank to stabilise the economy (i.e. the Taylor principle is satisfied) he forecast output and inflation tomorrow as a linear function of exogenous shocks today. That is, he forms expectations using the same functional form as under rational expectations, which he estimates using adaptive learning (Evans and Honkapohja, 2001). Conversely, if he does not trust the central bank to stabilise the economy he forecasts in output and inflation using a random walk. Every period he evaluates the forecasting performance of the two forecasting functions, adjusts his beliefs about the central bank's responsiveness, and changes the way he forms expectations going forward. In other words, he chooses between two forecasting heuristics in the sense of Brock and Hommes (1997) (so that the model henceforth will be called heuristic switching model). For example, if he observes that the random walk forecast yields smaller forecast errors, he perceives it to be more likely that the central bank is not committed to stabilise inflation. Therefore, he will put more weight on the random walk forecasting function going forward. This weight has a direct interpretation as the degree of de-anchoring because determines the pass-through of short-term conditions to long-term expectations. Thus, de-anchoring occurs when the central bank is not believed to stabilise inflation.

Figure 1: 10-year Ahead U.S. Inflation Expectations & Realised Inflation

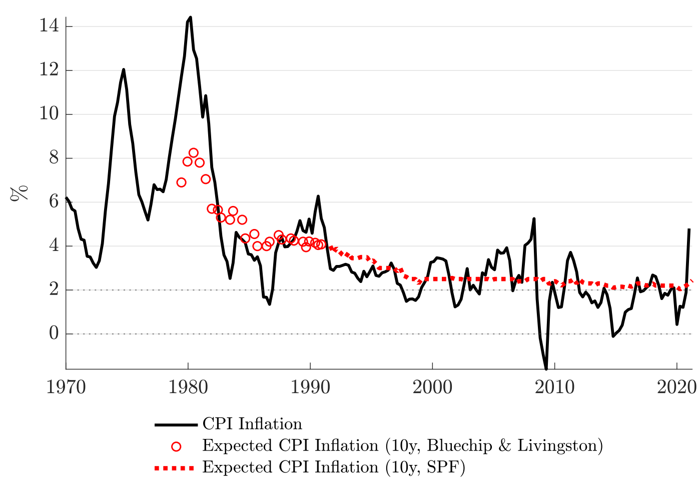

This theory of expectation formation embedded in an otherwise standard New Keynesian (NK) framework provides a novel, quantitatively realistic model of endogenous expectation de-anchoring. The model features the same steady state as under rational expectations (RE) but when expectations de-anchor, the volatility of output growth and inflation increases. As long as monetary policy is unconstrained, it can prevent expectations de-anchoring to cause inflationary or deflationary spirals as long as it reacts more than one-to-one to inflation, i.e. if the Taylor principle is satisfied. However, if monetary policy is constrained by the zero lower bound, de-anchoring can lead to self-reinforcing deflationary spirals (a feature absent from the standard rational expectations model and different from the self-fulfilling liquidity trap of Benhabib et al. (2002), for example). Thus, the potential welfare loss of de-anchoring is asymmetric and bigger in a low-interest rate environment. My results have implications regarding the monetary policy regime most likely to ensure macroeconomic stability. I find that Inflation Targeting ensures stationary model dynamics even when expectations are fully de-anchored. However, under Price Level Targeting there exists a threshold level of de-anchoring, above which the model becomes unstable unless the central bank is extremely responsive to inflation. This is the case because the combination of random walk forecasting and a history-dependent policy rule leads to explosive dynamics. To prevent this scenario, a central bank would have to react to inflation movements by adjusting the interest rate more than threefold for realistic slopes of the Phillips curve. Thus, my findings provide an even stronger case against PLT when agents are not fully rational than the arguments brought forward by Mele et al. (2020), contradicting the well-established result under rational expectations of Evans (2012), for example.

I subsequently estimate the model using a particle filter and U.S. expectations data. Here, I first establish that indeed, the model with forecast switching fits the data much better than a standard rational expectations model. In fact, this is a common property of models that deviate from rational expectations (e.g. Milani, 2007). Furthermore, the model-implied long-term inflation expectations fit the data quite well even though long-term expectations were not used in the estimation. Consequently, the model-implied long-term expectations the same degree of time-varying sensitivity to news as the observed forecasts of the SPF. Finally, the model-implied weight put on the random walk forecasting function implies that expectations were highly de-anchored in the early 1980s. Since the late 1980s expectations have been largely anchored, with temporary upticks in degree of de-anchoring coinciding with the recessionary periods of the early 1990s and 2000s. Since 2021, however, the estimated model detects a significant uptick in the de-anchoring of expectations.

After having established that the estimated model provides a realistic description of expectation formation, I use simulation methods to investigate the conduct of monetary policy when the anchoring of expectations is evolving endogenously. First, I show that the optimal inflation response in a Taylor rule significantly varies with the prevailing equilibrium interest rate: In a high nominal interest rate the central bank optimally responds less to inflation compared to the rational expectations model because overly aggressive interest rate fluctuations can lead to expectations de-anchoring themselves. That is, a central bank faces the trade-off between missing its inflation target, which can give rise to expectations de-anchoring, and excessive interest rate fluctuations which, depending on the expectation formation, can equally give rise to expectations de-anchoring (mirroring the results of Eusepi et al. (2020a) and Gáti (2023)). In a low nominal interest rate environment, i.e. when the Zero Lower Bound is often a binding constraint, the central bank optimally responds much more strongly to inflation compared to the rational expectations model, indicating how large the welfare loss of de-anchoring at the zero lower bound is. Second, I augment the analytical discussion of the stability properties of Inflation Targeting as well as Price Level Targeting by evaluating their respective welfare implications in the simulated model. Here, I show that in a high interest rate environment Price Level Targeting leads to higher welfare than Inflation Targeting, consistent with the literature. However, at the Zero Lower bound, Price level Targeting can be destabilising: Due to the zero lower bound it fails to generate to generate offsetting inflation necessary to make up for past inflation misses, which causes a rapid credibility loss and prevalence of random walk forecasting. This shift is much more detrimental under Price Level Targeting than under Price Level. In a final step of this simulation exercise, I confirm the optimal policy recommendation of Gáti (2023): a central bank can improve welfare by reacting forcefully to movements in long-term inflation expectations. This allows the central bank accommodate inflation fluctuations when expectations are well-anchored, while implicitly conditioning its reaction function on the degree of de-anchoring. However, my simulated model suggests that the central bank needs to be much less aggressive to achieve this goal than suggested by Gáti (2023).

With this set of results, I contribute to the literature on the conduct of monetary policy under endogenously (de-)anchored inflation expectations. Here, several studies have drawn on the seminal work of Marcet and Nicolini (2003) and modelled the anchoring of expectations by endogenising the learning gain in a standard adaptive learning set-up. In this approach, the learning gains is a function of realised forecast errors, resulting in an endogenously time-varying sensitivity of expectations to forecast errors. Using this approach, Gáti (2023) argues that monetary policy should respond more aggressively when expectations de-anchor to suppress the volatility caused by high degrees of de-anchoring. Even though this is a common result in the adaptive learning literature, in the context of de-anchoring it critically depends on the modelling of interest rate expectations: If interest rate expectations are subject to the same type of behavioural modelling, the central bank optimally responds less strongly to inflation (compared to rational expectations) because it otherwise adds additional volatility on top of the volatility coming from inflation expectations (Eusepi et al., 2020b; Gáti, 2023).

In contrast, I model de-anchoring using the endogenous switching between different forecasting functions: de-anchoring is not determined by the size of the learning gain, but depends on whether or not the central bank is trusted to stabilise the economy. Despite this different modelling set-up, there exists an implicit mapping between the two approaches: Modelling de-anchored expectations as a random walk forecast results in a series of aggregate expectations that display a realistic time-varying sensitivity to forecast errors, i.e. a time-varying learning gain. However, this setup of endogenous forecast switching allows me to analytically analyse the model stability for different degrees of de-anchoring under different policy regimes.

With this approach, my paper is also closely related to the literature on the conduct of monetary policy under endogenous central bank credibility that builds on the framework by Brock and Hommes (1997). Here, for example, Honkapohja and Mitra (2020) discuss the stability conditions of price level targeting. Further studies in this literature investigate, for example, forward guidance under endogenous credibility (Goy et al., 2020). In a similar fashion, Lansing (2021) models a rational agent who optimally updates his beliefs about the probability of being on the transition path to the targeted equilibrium or deflation equilibrium.

Against the backdrop of this literature, the first main contribution of my paper is that it provides a (potentially first[4]) New Keynesian model with endogenous expectations de-anchoring that is fully estimated using likelihood methods and expectations data. As a second main contribution, I use this model to discuss the conduct of monetary policy when the anchoring of expectations is determined endogenously. Here, my results demonstrate that it is crucial to distinguish between the threat of de-anchoring at the Zero Lower Bound from de-anchoring away from the ZLB, as the two cases carry drastically different outcomes and call for different policy measures.

More generally, this paper is connected to the wide literature on the empirical properties of inflation expectations, especially the part of the literature investigates whether expectations are anchored or not (e.g. Levin et al., 2004; Beechey et al., 2011; Strohsal et al., 2016; Reis, 2021). Here, Carvalho et al. (2023) estimate a model of expectations with endogenous gain learning in a partial equilibrium set-up. Their estimates indicate that expectations were poorly anchored before the late 1990s, with long-run expectations displaying high sensitivity

[4]To the best of the author's knowledge this is the first paper to do so. This stands in contrast to Carvalho et al. (2023), who estimate the process for inflation as well as inflation expectations only; to Gáti (2023), who uses simulated method of moments to estimate a subset of parameters; and to Ozden (2021) & Ozden and Wouters (2021) who do not use expectations data to estimate the parameters of their endogenous expectations switching model.

to new information. In this context, the main contribution of this paper is that the estimated model yields an extended measure expectations de-anchoring. This measure suggests that expectations became de-anchored to a significant degree starting in 2021.

The paper proceeds as follows. Section 2 introduces the model with heterogeneous expectations. Section 3 describes the model solution and stability properties. Section 4 describes the estimation procedure and illustrates the dynamics of the estimated model. Section 5.2 analyses the performance of different monetary policy rules by simulating the estimated model. Section 6 concludes.

2. A Model of De-anchoring

I use a standard New Keynesian (NK) model to study the de-anchoring of expectations. To isolate the effect of de-anchoring, I keep the model simple and only introduce one distortion, sticky prices à la Rotemberg (1982). In the following, I first lay out the non-linear version of the model before turning to the linearised version and discussing the expectation formation.

2.1. The New Keynesian Framework

The model is populated by a representative household, intermediate and final goods producers, a central bank, and the government. The representative chooses consumption C_t, labour H_t, and government bonds B_t to maximise the expected discounted stream of utility

$$\hat{\mathbb{E}}_0 \sum_{t=1}^{\infty} \beta^t \Xi_{t-1} \left[\frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{H_t^{1+\eta}}{1+\eta} \right] \tag{1}$$

subject to the budget constraint

$$P_t C_t + B_t = P_t W_t H_t + R_{t-1} B_{t-1} + P_t Div_t + T_t \tag{2}$$

where $\hat{\mathbb{E}}$ is the non-rational expectations operator, Ξ_t is a shock to the discount factor, P_t is the price level, W_t is the real wage, R_t is the gross interest rate, T_t is a lump-sum tax and Div_t are real profits from the intermediate good firms. B_t denotes the one-period government bonds in zero net supply. Solving the representative household's problem yields the intertemporal Euler equation

$$1 = \beta R_t \hat{\mathbb{E}}_t \frac{\Xi_t}{\Xi_{t-1}} \left(\frac{C_t}{C_{t+1}} \right)^{\sigma} \frac{1}{\Pi_{t+1}} \tag{3}$$

and the intratemporal labour supply condition

$$W_t = \chi H_t^{\eta} C_t^{\sigma} \tag{4}$$

where $\zeta_t = \frac{\Xi_t}{\Xi_{t-1}}$ follows an AR(1) process in logs.

The **final goods producers** transform intermediate goods into the homogeneous good, which is obtained by aggregating intermediate goods using the following technology

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\epsilon-1}{\epsilon}} dj \right)^{\frac{\epsilon}{\epsilon-1}} \tag{5}$$

where $Y_t(i)$ is the intermediate good of firm i. The price index for the aggregate homogeneous good is

$$P_t = \left[\int_0^1 P_t(i)^{1-\epsilon} dj \right]^{\frac{1}{1-\epsilon}} \tag{6}$$

and the demand for the differentiated good $j \in (0, 1)$ is

$$Y_t(j) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon} Y_t \tag{7}$$

The **intermediate goods producer** i produces using labour as the only input with a constant return to scales technology

$$Y_t(i) = AH_t \tag{8}$$

where A denotes the total factor productivity. The firm i sets the price $P_t(i)$ of its differentiated good to maximise its profits

$$\max_{P_{t+j}(i)} E_t \sum_{j=0}^{\infty} \Lambda_{t,t+j} \left(\left(\frac{P_{t+j}(i)}{P_{t+j}} - \frac{MC_{t+j}}{P_{t+j}} \mathcal{M}_{t+j} \right) Y_{t+j}(i) - \frac{\varphi}{2} \left(\frac{P_{t+j}(i)}{P_{t+j-1}(i)} - 1 \right)^2 Y_{t+j} \right) \tag{9}$$

subject to the demand curve for intermediate goods and where $\Pi \geq 1$ is the steady-state inflation rate, $\mathcal{M}_t$ is an exogenous mark-up shock that follows an AR(1) process in logs, and $\Lambda_{t,t+1}$ is the household's stochastic discount factor. The parameter $\varphi \geq 0$ measures the cost of price adjustment in units of the final good with $\varphi = 0$ leading to the flexible-price output level Y_t^*. Solving the intermediates firm problem, the optimal pricing rule implies in a symmetric equilibrium that

$$\varphi \left(\frac{\Pi_t}{\bar{\Pi}} - 1 \right) \frac{\Pi_t}{\bar{\Pi}} = (1 - \epsilon) + \epsilon MC_t \mathcal{M}_t + \varphi \hat{\mathbb{E}}_t \Lambda_{t,t+1} \left(\frac{\Pi_{t+1}}{\bar{\Pi}} - 1 \right) \frac{\Pi_{t+1} Y_{t+1}}{\bar{\Pi} Y_t} \tag{10}$$

The fiscal authority sets taxes to balance the budget in every period so that the **aggregate resource constraint** is

$$C_t = Y_t \left[1 - \frac{\varphi}{2} \left(\frac{\Pi_t}{\bar{\Pi}} - 1 \right)^2 \right] \tag{11}$$

Combining and linearising Equations (1) - (11) around the zero-inflation steady state yields the familiar 3 equation system (where lower case letters symbolise deviations of the respective variable from its steady state)

$$
\begin{aligned}
y_t &= \hat{\mathbb{E}}_t \left\{y_{t+1}\right\} - \tfrac{1}{\sigma}\Big(r_t - \hat{\mathbb{E}}_t \pi_{t+1} + \xi_t\Big) \\
\pi_t &= \beta \hat{\mathbb{E}}_t \left\{\pi_{t+1}\right\} + \underbrace{\frac{\varepsilon - 1}{\varphi}\left(\sigma + \chi\eta\right)}_{\equiv\kappa} y_t + \underbrace{\frac{\varepsilon - 1}{\varphi}}_{\equiv\lambda} \mu_t
\end{aligned} \tag{12}
$$

with

$$
\begin{aligned}
\mu_t &= \rho_a \mu_{t-1} + \varepsilon_t^{\mu} \\
\xi_t &= \rho_\xi \xi_{t-1} + \varepsilon_t^{\xi}
\end{aligned} \tag{13}
$$

where μ_t is the log-linearised mark-up shock. The shocks ε_t^{μ} and ε_t^{ξ} are distributed normally with standard deviations σ_μ and σ_ξ around a mean of zero.

To close the model, I consider two different monetary policy frameworks for the **central bank**, Inflation Targeting and Price Level Targeting. Under Inflation Targeting, the central bank follows a standard Taylor rule subject to the zero lower bound

$$
r_t = \max\left[0, \bar{\pi} + \phi_\pi(\pi_t - \bar{\pi}) + \phi_y y_t + m_t\right] \tag{14}
$$

where m_t evolves according to

$$
m_t = \rho_m m_{t-1} + \varepsilon_t^r \quad \varepsilon_t^r \sim \mathcal{N}(0,\, \sigma_r)
$$

Under Price Level Targeting the central bank reacts to the deviation of the price level p_t from the target price level path p_t^*

$$
r_t = \max\left[0, \bar{\pi} + \phi_\pi \hat{p}_t + \phi_y y_t + m_t\right] \tag{15}
$$

where, for some given π_0, the price level gap $\hat{p}_t$ evolves according to $\hat{p}_t = \pi_t - \bar{\pi} + \hat{p}_{t-1}$ and $\hat{p}_{t-1}$ enters the model as an additional state variable.

2.2. Endogenous Anchoring of Expectations

To model the endogenous (de-)anchoring of expectations, I deviate from rational expectations and assume that the representative household does not know the true structure of the economy. Instead, he forms believes about the state of the world he is in. That is, he weighs his beliefs about whether or not the Taylor principle is satisfied: if he trusts the central bank to stabilise the economy (i.e. he believes the Taylor principle to be satisfied) he forecast output and inflation tomorrow as a linear function of the exogenous states today. That is, he forms expectations using the same functional form as under rational expectations, which he estimates using adaptive learning (Evans and Honkapohja, 2001). Conversely, if he does not trust the central bank to stabilise the economy he forecasts output and inflation using a random walk. Every period he evaluates the forecasting performance of the two heuristics, adjusts his beliefs about the central bank's responsiveness, and changes the way he forms expectations going forward. That is, he chooses between two forecasting heuristics in the sense of Brock and Hommes (1997). The agent's average forecast is a weighted average of the two forecasting heuristics, i.e.

$$\hat{\mathbb{E}}_t \mathbf{z}_{t+1} = (1 - n_t)\hat{\mathbb{E}}_t^a \mathbf{z}_{t+1} + n_t \hat{\mathbb{E}}_t^n \mathbf{z}_{t+1} \tag{16}$$

where $\hat{\mathbb{E}}^a$ denotes the expectations under the adaptive learning heuristic, $\hat{\mathbb{E}}^n$ denotes the expectations under the naive random walk heuristic, n_t is the weight put on the random walk forecasting function, and $\mathbf{z}_t$ is a vector containing the endogenous realisation $y_t, \pi_t,\, r_t$ (and, in the case of Price Level Targeting, also the price level gap $\hat{p}_t$).

If the household believes that the central bank stabilises the economy, he assumes that the endogenous variables $\mathbf{z}_t$ evolve according to the following linear perceived law of motion (PLM):

$$\mathbf{z}_t = \Psi_t \mathbf{x}_{t-1} + \epsilon_t \tag{17}$$

where $\boldsymbol{\Psi}_t = [\mathbf{a}_t,\, \mathbf{b}_t]$ and $\mathbf{x}_t = (1,\, \xi_t,\, \mu_t\, m_t)'$ so that the corresponding forecast becomes

$$\hat{\mathbb{E}}_t^a \mathbf{z}_{t+1} = \Psi \mathbf{x}_t \tag{18}$$

Under Price Level Targeting the adaptive forecasting heuristic additionally makes use of the lagged price level gap to forecast $\mathbf{z}$ so that $\boldsymbol{\Psi}_t = [\mathbf{a}_t,\, \mathbf{b}_t,\, \tilde{\mathbf{c}}_t]$ and $\mathbf{x}_t = (1,\, \xi_t,\, \mu_t\, m_t,\, \hat{p}_{t-1})'$.[5] In both cases, the representative household forecasts inflation and output tomorrow as a function of the exogenous states of the economy today, as under rational expectations. However, he does not know the exact quantitative relationship between the exogenous shocks and the endogenous variables. Therefore, each period the household re-estimates the perceived law of motion and updates the parameters $\boldsymbol{\Psi}_t$ following:

$$\begin{gathered} \boldsymbol{\Psi}_t' = \boldsymbol{\Psi}_{t-1}' + \gamma \boldsymbol{R}_{t-1}^{-1} \boldsymbol{x}_{t-2} \left(\boldsymbol{z}_{t-1} - \boldsymbol{\Psi}_{t-1} \boldsymbol{x}_{t-2} \right)' \\ \mathbf{R}_t = \mathbf{R}_{t-1} + \gamma \left(\mathbf{x}_{t-2} \mathbf{x}_{t-2}' - \mathbf{R}_{t-1} \right) \end{gathered} \tag{19}$$

where γ is the constant gain parameter which determines the weight given to new information. A constant γ implies that recent observations receive a higher weight than older observations (resembling the experience-based learning documented by Malmendier and Nagel (2016)). If instead $\gamma = 1/t$, each observation would be weighted equally, and the updating of

[5] $\tilde{\mathbf{c}}_t$ is a 3×1 vector of coefficients assigned to the lagged price level gap and denoted by a tilde because later it will be convenient to define $\mathbf{c} = [\mathbf{0}_{3\times 2}\ \tilde{\mathbf{c}}]$

coefficients would be equivalent to running a recursive OLS regression in every period.

If the household does not believe that the central bank will stabilise the economy, he forecasts that output and inflation follow a random walk (i.e. he believes that $\phi_\pi = 1$):

$$\hat{\mathbb{E}}_t^n \mathbf{z}_{t+1} = \hat{\mathbb{E}}_t^n \mathbf{z}_t = \mathbf{z}_{t-1} \tag{20}$$

Each period, the household re-evaluates the likelihood of being in each of the two worlds, i.e. whether or not the Taylor Principle is satisfied. He does so by evaluating the forecast performance of the two forecasting functions in terms of the sum of squared forecast errors. A logistic function maps the forecast errors into a weight n_t that represents the likelihood of being in a world with passive monetary policy:[6]

$$n_t = \frac{\exp\left(-\theta e_{t-1}^n\right)}{\exp\left(-\theta e_{t-1}^n\right) + \exp\left(-\theta e_{t-1}^a\right)} \tag{21}$$

with

$$e_t^i = (1-\omega) e_{t-1}^i + \omega \varepsilon_t^i \quad i \in \{a, n\} \tag{22}$$

where e_{t-1}^i is the recursive weighted average of the sum of squared forecast errors ε_t^i. That is, the household avoids making systematic forecasting mistakes by switching to the better performing forecast rule. The two parameters ω and θ jointly determine how the household interprets the performance of the two forecasting functions: ω controls the rate at which previous forecast errors decay, i.e. the length of households' memory. θ determines the sensitivity to differences in the forecasting performance: when $\theta \to \infty$, the household immediately switches to the forecasting rule with the lower forecast error. Conversely, with $\theta = 0$ the household never switches and the weights remain fixed at 0.5. Any intermediate value implies an imperfect adoption of the better performing forecast rule. In the steady state the weight put on each forecasting function is equal to $n = 0.5$ because there are no exogenous shocks and neither forecasting function makes any forecast error.

This heuristic switching mechanism provides an easy statistic measuring the degree of expectation anchoring: when expectations are well anchored, the adaptive rule will produce good forecasts and the weight put on the random walk forecasting statistics will go down. When expectations are not well anchored, for example, because the central bank does not react enough to offset shocks that move away output and inflation from their target, the adaptive rule will produce bad forecasts and more weight will be put on the random walk forecasting heuristic. Monetary policy can affect these dynamics by sending correct signals for the evolutionary selection of the heuristics and induce stable dynamics converging to the

[6] This expression naturally arises if the forecasting performance is not perfectly observed (a reasonable assumption given the frequent revisions of gdp-related measures) but instead is observed with a logistically distributed noise term (e.g. Anufriev et al., 2013; De Grauwe and Ji, 2019).

rational expectations steady state.

2.3. Implied Long-term Expectations

The heuristic switching model implicitly also pins down long-term expectations. The long-term forecast (i.e. the forecast for a horizon k sufficiently large) of the adaptive heuristic is simply the intercept of the learning rule:

$$\hat{\mathbb{E}}_t^a \mathbf{z}_{t+k} = \mathbf{a}_t \tag{23}$$

Under the naive forecasting heuristic, the household extends his forecasts over horizon k as the recursive weighted average of realisations over the previous k periods.

$$\hat{\mathbb{E}}_t^n \mathbf{z}_{t+k} = (1 - {}^1\!/\!_k)\hat{\mathbb{E}}_{t-1}^n \mathbf{z}_{t+k-1} + {}^1\!/\!_k\, \mathbf{z}_{t-1} \tag{24}$$

For $k = 1$ this nests the one quarter ahead forecast discussed previously. Furthermore, for $k \to \infty$ the naive forecast will converge to the mean of the process $\mathbf{z}$, which corresponds to the long-run forecast of the adaptive heuristic.

The average long-run forecast is a weighted average of the two heuristics. As before, the weight put on the random walk forecast is pinned down by the accuracy of the one-quarter ahead forecast.

$$\hat{\mathbb{E}}_t \mathbf{z}_{t+k} = (1 - n_t)\hat{\mathbb{E}}_t^a \mathbf{z}_{t+k} + n_t \hat{\mathbb{E}}_t^n \mathbf{z}_{t+k} \tag{25}$$

$$= (1 - n_t)\mathbf{a} + n_t((1 - {}^1\!/\!_k)\hat{\mathbb{E}}_{t-1}^n \mathbf{z}_{t+k-1} + {}^1\!/\!_k\, \mathbf{z}_{t-1}) \tag{26}$$

Therefore, the sensitivity of long-term inflation expectations to short-term developments increases if the household doubts the central bank's ability to stabilise inflation, because it leads to a higher weight put on recent inflation developments. This expectation formation process is empirically plausible: following a monetary policy shock, these expectations react in the same way as observed long-run expectations (see Appendix C.1).

3. Model Equilibria and Stability

In this section I solve the model under the two different monetary policy frameworks, Inflation Targeting and Price Level Targeting, and discuss the stability properties both away from and at the ZLB. In each case, I first derive the fixed points of the adaptive forecasting heuristic under some fixed weight n put on the random walk forecasting function. In a second step, I use these fixed points to characterise the steady state under time-varying n_t. In a third step, I discuss the E-Stability of this steady state, i.e. whether the steady state can be learned.[7]

[7] E-stability ensures that estimated coefficients of the learners converge to their respective fixed points.

3.1. Dynamics under Inflation Targeting

To simplify the analysis, I assume that the central bank only reacts to deviations of inflation from target (i.e. $\phi_y = 0$). Furthermore, I initially consider an environment without the zero lower bound in 3.1.1 before reintroducing it in 3.1.2.

3.1.1. Away from the ZLB

Substituting the monetary policy rule as well as the perceived laws of motion (17) & (20) into the system of equations (12) and rearranging yields the Actual Law of Motion (ALM)

$$\begin{aligned} y_t &= (1-n_t)[a_y + b_{y,\xi}\xi_t + b_{y,\mu}\mu_t + b_{y,m}m_t] + n_t y_{t-1} - \frac{1}{\sigma}(\bar{\pi} + \phi_\pi(\pi_t - \bar{\pi}) \\ &\quad + m_t - (1-n_t)[a_\pi + b_{\pi,\xi}\xi_t + b_{\pi,\mu}\mu_t + b_{\pi,m}m_t] - n_t y_{t-1} + \xi_t) \\ \pi_t &= \beta[(1-n_t)[a_\pi + b_{\pi,\xi}\xi_t + b_{\pi,\mu}\mu_t + b_{\pi,m}m_t] + n_t\pi_{t-1}] + \kappa y_t + \lambda\mu_t \end{aligned} \tag{27}$$

or in matrix notation

$$\Leftrightarrow \mathbf{z}_t = \mathbf{A}\underbrace{[(1-n_t)(\mathbf{a} + \mathbf{b}\mathbf{w}_t) + n_t\mathbf{z}_{t-1}]}_{\hat{\mathbb{E}}_t \boldsymbol{z}_{t+1}} + \mathbf{B}\mathbf{w}_t + \mathbf{C}\bar{\mathbf{z}} \tag{28}$$

Substituting (28) into the ODE associated with the updating equations (19) yields

$$\frac{\partial \boldsymbol{\Psi}'}{\partial \tau} = \mathbf{R}^{-1}E\mathbf{x}_{t-2}\left(\mathbf{A}\left[(1-n)\boldsymbol{\Psi}\mathbf{x}_{t-1})\right] + \mathbf{A}n\mathbf{z}_{t-2} + \mathbf{B}\mathbf{w}_t + \mathbf{C}\bar{\mathbf{z}}\right) - \boldsymbol{\Psi}\mathbf{x}_{t-2})' \tag{29}$$

where τ denotes *notional time* (Evans and Honkapohja, 2001). As long as the ALM in Equation (28) is asymptotically stationary, i.e. has roots within the unit circle so that it is mean-reverting

$$\lambda < 1 \quad \forall \lambda \in \Lambda = \{\Lambda : |\mathbf{I} - n\mathbf{A} - \boldsymbol{\Lambda}\mathbf{I}| = 0\} \tag{30}$$

and when holding n fixed, the coefficients of the adaptive PLM converge to their respective fixed points (for a detailed derivation see Appendix A)

$$vec(\bar{\boldsymbol{a}}') = (\boldsymbol{I} - \boldsymbol{A})^{-1} vec\left((\boldsymbol{C}\bar{\boldsymbol{z}})'\right) \tag{31}$$

$$vec(\bar{\boldsymbol{b}}') = (\boldsymbol{I} - \mathbf{G}_1)^{-1}\mathbf{G}_2 vec(\mathbf{B}') \tag{32}$$

where

$$\begin{aligned} \mathbf{G}_1 &= \mathbf{A}(1-n) \otimes \mathbf{F} + \mathbf{A}n \otimes \boldsymbol{\Sigma}_{\mathbf{w}}{}^{-1}\left(\boldsymbol{I} - n\mathbf{A} \otimes \mathbf{F}\right)\mathbf{A}(1-n) \otimes \boldsymbol{\Sigma}_{\mathbf{w}} \\ \mathbf{G}_2 &= \mathbf{I} \otimes \mathbf{F} + \mathbf{A}n \otimes \boldsymbol{\Sigma}_{\mathbf{w}}{}^{-1}\mathbf{I} \otimes \boldsymbol{\Sigma}_{\mathbf{w}}\left(\boldsymbol{I} - n\mathbf{A} \otimes \mathbf{F}\right)\mathbf{I} \otimes \boldsymbol{\Sigma}_{\mathbf{w}} \end{aligned}$$

Holding n fixed, the mapping T from the perceived to the actual law of motion can therefore be characterised as

$$T(\mathbf{\Psi}', n) = \begin{pmatrix} T(\mathbf{a}', n) \\ T(\mathbf{b}', n) \end{pmatrix} = \begin{pmatrix} \left((\boldsymbol{I} - \boldsymbol{A}n)^{-1}\boldsymbol{A}(1-n)\right)' \\ \mathbf{G}_1 \end{pmatrix} \tag{33}$$

This result is summarised in the following proposition:

Proposition 1. *The mapping (33) from the PLM to the ALM has a unique fixed point* $\forall n \in [0, 1]$ *if the Taylor principle is satisfied by the central bank, i.e.*

$$\phi_\pi > 1$$

Proof. See Appendix A.1. □

The fixed point $\bar{\mathbf{a}}$ is independent of n and equivalent to the steady state under rational expectations:

$$\bar{\mathbf{a}} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{C}\bar{\mathbf{z}} = \begin{pmatrix} 0 & \frac{1-\beta}{\kappa} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \bar{y} \\ \bar{\pi} \end{pmatrix} \tag{34}$$

On the other hand, $\bar{\mathbf{b}}$ is a non-linear function in n. Intuitively, two offsetting effects determine the magnitude of $\bar{\mathbf{b}}$: as n decreases $\mathbf{w}_t$ has a stronger effect on $\mathbf{z}_t$ because the average forecast of $\mathbf{z}_{t+1}$ loads stronger on $\mathbf{w}_t$. At the same time, as n decreases, the correlation between $\mathbf{z}_{t+1}$ and $\mathbf{z}_t$ decreases so that $\mathbf{w}_t$ has a smaller effect on $\mathbf{z}_{t+1}$.

Now, if n_t is allowed to vary over time, only $n_t = 0.5$ is consistent with a steady state in which both heuristics do not make any forecast error. This allows me to characterise a unique steady state under heterogeneous expectations:

Proposition 2. *The dynamic system (27) has a unique steady state with* $\pi^* = \bar{\pi}$*,* $y^* = \frac{1-\beta}{\kappa}\bar{\pi}$*,* $r^* = \bar{\pi}$ *and* $n^* = 0.5$

Proof. See Appendix A.2 □

For this steady state to be stable it is necessary that 1) the process $\mathbf{z}_t$ is stationary for $n = 0.5$ so that the fixed point of the mapping T exists and that 2) the steady state is e-stable, i.e. learnable. According to the E-stability principle (Evans and Honkapohja, 2001, Chapter 13), convergence under learning requires that all eigenvalues of

$$DT_{\mathbf{\Psi}}(\mathbf{\Psi}', n)) - \mathbf{I} \tag{35}$$

need to have negative real parts to ensure that estimated coefficients of adaptive forecasting function converge to their respective fixed points. Holding n^* fixed at the steady state and assuming that the Taylor principle is satisfied (see Proposition 1), the stationarity and E-stability requirements are always satisfied for the steady state away from the ZLB:

Proposition 3. *The steady state away from the zero lower bound is stable.*

Proof. See Appendix A.3. □

Figure 2: Stability Conditions

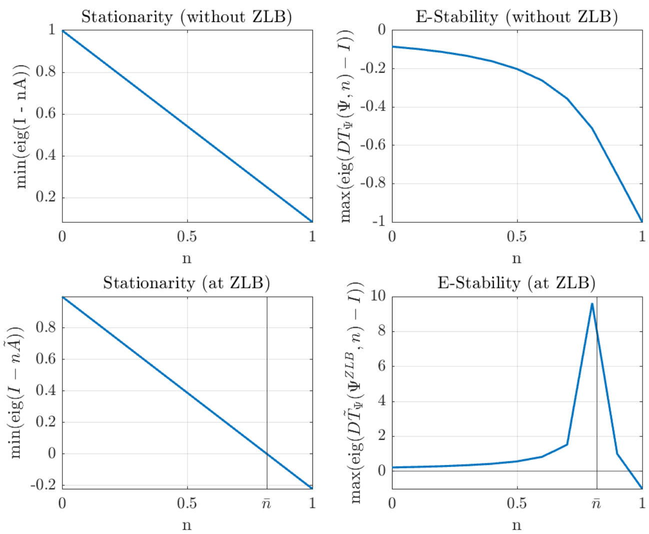

Note: This figure shows the stability properties of the model under Inflation Targeting as a function of the weight n put on the random walk forecasting function without and with the ZLB. For stationarity, the *minimum* eigenvalue needs to be positive. For E-stability the *maximum* eigenvalue needs to be negative.

Away from the steady state, the forecast switching can create, or at least amplify, booms and busts. For example, if the adaptive forecast function underestimate inflation over several quarters, the resulting deterioration in forecasting performance will increase the weight of the random walk forecasting function. The resulting drift of inflation and output decreases the weight of the minimum state variable forecast. This de-anchoring of expectations can lead to self-reinforcing cycles which amplify the drift and poses additional challenges to monetary policy not present in the standard rational expectations framework.

3.1.2. At the ZLB

Re-introducing the ZLB into the system, the ALM at the ZLB in matrix notation becomes

$$\mathbf{z}_t = \tilde{\mathbf{A}}\left[(1-n_t)(\mathbf{a}+\mathbf{b}\mathbf{w}_t) + n_t\mathbf{z}_{t-1}\right] + \tilde{\mathbf{B}}\mathbf{w}_t + \tilde{\mathbf{C}}\bar{\mathbf{z}} \tag{36}$$

where the matrices containing the coefficients of the ZLB ALM are denoted with a tilde. Following the same approach as before, it can be shown that the ALM has a fixed point at

$$vec(\tilde{\bar{\mathbf{a}}}) =(\boldsymbol{I} - \tilde{\boldsymbol{A}})^{-1} vec\left((\tilde{\boldsymbol{C}}\bar{\boldsymbol{z}})'\right) \tag{37}$$

$$vec(\tilde{\bar{\mathbf{b}}}') = \left(\boldsymbol{I} - \tilde{\mathbf{G}}_1\right)^{-1} \tilde{\mathbf{G}}_2 vec(\tilde{\mathbf{B}}') \tag{38}$$

where

$$\tilde{\mathbf{G}}_1 = \tilde{\mathbf{A}}(1-n) \otimes \mathbf{F} + \tilde{\mathbf{A}}n \otimes \mathbf{\Sigma_w}^{-1} \left(\boldsymbol{I} - n\tilde{\mathbf{A}} \otimes \mathbf{F}\right) \tilde{\mathbf{A}}(1-n) \otimes \mathbf{\Sigma_w}$$

$$\tilde{\mathbf{G}}_2 = \mathbf{I} \otimes \mathbf{F} + \tilde{\mathbf{A}}n \otimes \mathbf{\Sigma_w}^{-1} \mathbf{I} \otimes \mathbf{\Sigma_w} \left(\boldsymbol{I} - n\tilde{\mathbf{A}} \otimes \mathbf{F}\right) \mathbf{I} \otimes \mathbf{\Sigma_w}$$

Holding n fixed, the mapping $\tilde{T}$ from the perceived to the actual law of motion for the PLM coefficients at the ZLB is now given by

$$\tilde{T}(\mathbf{\Psi}', n) = \begin{pmatrix} \tilde{T}(\mathbf{a}', n) \\ \tilde{T}(\mathbf{b}', n) \end{pmatrix} = \begin{pmatrix} \left((\boldsymbol{I} - \tilde{\boldsymbol{A}}n)^{-1} \tilde{\boldsymbol{A}}(1-n)\right)' \\ \tilde{\mathbf{G}}_1 \end{pmatrix} \tag{39}$$

However, this mapping has a fixed point only for certain values of n:

Proposition 4. *The mapping (39) from the PLM to the ALM has a unique fixed point for all* $n \in [0, \bar{n})$

Proof. See Appendix A.4 □

For all values of $n < \bar{n}$ the fixed point of $\tilde{\bar{\mathbf{a}}}$ exists and is independent of n. Furthermore, it is equivalent to the steady state under rational expectations:

$$\tilde{\bar{\mathbf{a}}} = (\mathbf{I} - \tilde{\mathbf{A}})^{-1} \tilde{\mathbf{C}} \bar{\mathbf{z}} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \bar{y} \\ \bar{\pi} \end{pmatrix} \tag{40}$$

Again, if n_t is allowed to vary over time, only $n_t = 0.5$ is consistent with a steady state in which neither heuristic makes any forecast error. This allows me to characterise a second steady state under heterogeneous expectation:

Proposition 5. *There exists a second steady state with* $\tilde{\pi}^* = 0$, $\tilde{y}^* = 0$, $\tilde{r}^* = 0$ *and* $n^* = 0.5$

Proof. See Appendix A.5 □

However, the ZLB steady state is not stable because either $\bar{n} < 0.5$ so that the process $\mathbf{z}_t$ is not stationary at the steady state or because it is not E-stable:

Proposition 6. *The zero lower bound steady state is either not asymptotically stationary or not E-stable.*

Proof. See Appendix A.6 □

In general, at the ZLB both conditions (asymptotic stationarity and E-stability) never hold at the same time. This is because the eigenvalues of $D\tilde{T}_{\mathbf{\Psi}}(\mathbf{\Psi}', n) - \mathbf{I}$ have negative real parts only if the eigenvalues of $n\tilde{\mathbf{A}}$ are outside the unit circle, as Figure 2 illustrates.

To summarise, under Inflation Targeting the model has two steady states: one steady state away from the ZLB and one at the ZLB, but only the former is stable. These steady states are identical to the ones under rational expectations. However, the model features complex dynamics away from the steady state: whenever the households puts more (or less) weight on the random walk forecasting function, the relationship between exogenous and endogenous realisations changes so that adaptive learning rules will change. This in turn affects performance of the random walk forecasting function, creating fluctuations around the steady state.

3.2. *Dynamics under Price Level Targeting*

Under Price Level Targeting the fixed points $\bar{\mathbf{a}}, \bar{\mathbf{b}}, \bar{\mathbf{c}}$ of the adaptive forecasting heuristic are all (non-linearly) dependent on each other so that no closed-form solution can be derived. To characterise the solution at least computationally, I keep the assumption of $\phi_y = 0$ and, furthermore, assume that the shocks $\mathbf{w}_t$ are iid. Under the adaptive forecasting function, the household, therefore, only forecasts the mean values of the state variables, i.e. engages in steady-state learning. This does not change the stability properties of the system but simplifies the analysis. As before, I first consider the steady state in an environment without the zero lower bound before reintroducing it.

3.2.1. *Away from the ZLB*

Substituting the monetary policy rule (15) as well as the perceived laws of motion (20) & (17) into the system of equations (12) and rearranging yields the ALM

$$
\begin{aligned}
y_t &= (1-n_t)[a_y + b_{y,\xi}\xi_t + b_{y,\mu}\mu_t + b_{y,m}m_t + c_{y,\hat{p}}\hat{p}_{t-1}] + n_t y_{t-1} - \frac{1}{\sigma}(\bar{\pi} + \phi_\pi(\pi_t - \bar{\pi} + \hat{p}_{t-1}) \\
&\quad + m_t - (1-n_t)[a_\pi + b_{\pi,\xi}\xi_t + b_{\pi,\mu}\mu_t + b_{\pi,m}m_t + c_{\pi,\hat{p}}\hat{p}_{t-1}] - n_t y_{t-1} + \xi_t) \\
\pi_t &= \beta[(1-n_t)[a_\pi + b_{\pi,\xi}\xi_t + b_{\pi,\mu}\mu_t + b_{\pi,m}m_t + c_{\pi,\hat{p}}\hat{p}_{t-1}] + n_t\pi_{t-1}] + \kappa y_t + \lambda\mu_t \\
\hat{p}_t &= \pi_t - \bar{\pi} + \hat{p}_{t-1}
\end{aligned}
\tag{41}
$$

Or in matrix notation:

$$
\boldsymbol{z}_t = \boldsymbol{A}[(1-n_t)(\boldsymbol{a} + \boldsymbol{b}\boldsymbol{w}_t + \mathbf{c}\mathbf{z}_{t-1}) + n_t\boldsymbol{z}_{t-1}] + \boldsymbol{D}\boldsymbol{z}_{t-1} + \boldsymbol{B}\boldsymbol{w}_t + \boldsymbol{C}\bar{\boldsymbol{z}} \tag{42}
$$

where $\mathbf{c} = [\mathbf{0}_{3\times 2}\ \tilde{\mathbf{c}}]$. Substituting this into the ODE of the updating equations (19) yields[8]

$$\frac{\partial \mathbf{\Psi}'}{\partial \tau} = \mathbf{R}^{-1} E\mathbf{x}_{t-2}\left(\mathbf{A}\left[(1-n)(\mathbf{a}+\mathbf{b}\mathbf{w}_t+\mathbf{c}\mathbf{z}_{t-2})\right]+\mathbf{A}n\mathbf{z}_{t-2}+\mathbf{B}\mathbf{w}_t+\mathbf{C}\bar{\mathbf{z}})-\mathbf{\Psi}\mathbf{x}_{t-2}\right)' \tag{43}$$

To characterise the fixed points of the adaptive forecasting heuristic, I start with the fixed point $\bar{\mathbf{c}}$ which determines the other two fixed points. Unfortunately, the associated mapping

$$\frac{\partial \mathbf{c}'}{\partial \tau} = \mathbf{\Sigma}_p^{-1}\mathbf{\Sigma}_{p_{t-2},p_{t-1}}^{-1}(\mathbf{A}(1-n)\mathbf{c}_{:,3}+\mathbf{D}_{:,3}+n\mathbf{A}_{:,3})+\mathbf{\Sigma}_p^{-1}\mathbf{\Sigma}_{p_{t-2},(y_{t-1},\pi_{t-1})}(\mathbf{D}_{1:2}+n\mathbf{A}_{1:2}) \tag{44}$$

does not allow a closed form solution of the fixed point and must be solved numerically for $\bar{\mathbf{c}}$. Having pinned down $\bar{\mathbf{c}}$ this way, we can write (holding $(\mathbf{\Psi}, n)$ fixed) $\lim_{t\to\infty} E\mathbf{z_t}$ as

$$\mathbf{z}_t(\mathbf{\Psi}, n) = \left((\mathbf{I}-\mathbf{A}(1-n)\mathbf{c}-n\mathbf{A}-\mathbf{D})^{-1}\left(\mathbf{A}\left[(1-n)\mathbf{a}\right]+\mathbf{C}\bar{\mathbf{z}}\right)\right)' \tag{45}$$

as long as the ALM in Equation (42) is asymptotically stationary, i.e.

$$\lambda < 1 \quad \forall \lambda \in \Lambda = \{\Lambda : |\mathbf{I}-\mathbf{A}(1-n)\mathbf{c}-n\mathbf{A}-\mathbf{D}-\mathbf{\Lambda}\mathbf{I}| = 0\} \tag{46}$$

The fixed point $\bar{\mathbf{a}}$ then becomes

$$vec(\bar{\boldsymbol{a}}') = \left[(\boldsymbol{I}+(\boldsymbol{A}(1-n)\bar{\mathbf{c}}+n\mathbf{A}+\mathbf{D})(\boldsymbol{I}-(\boldsymbol{A}(1-n)\bar{\mathbf{c}}-n\mathbf{A}-\mathbf{D})^{-1}-\mathbf{A}(1-n)\right]^{-1} vec\left((\boldsymbol{C}\bar{\boldsymbol{z}})'\right) \tag{47}$$

Finally, the fixed point $\bar{\mathbf{b}}$ is then given by

$$vec(\bar{\boldsymbol{b}}') = \left[\boldsymbol{I}-(\mathbf{A}(1-n)\bar{\mathbf{c}}+\mathbf{D}+n\mathbf{A})\otimes\mathbf{I}\,\mathbf{A}(1-n)\otimes\mathbf{I}\right]^{-1}(\mathbf{A}(1-n)\bar{\mathbf{c}}+\mathbf{D}+n\mathbf{A})\otimes\mathbf{I}\,vec(\mathbf{B})' \tag{48}$$

Holding n fixed, the mapping T from the perceived to the actual law of motion for the PLM coefficients is now given by

$$T(\mathbf{\Psi}', n) = \begin{pmatrix} ((\mathbf{I}+(\mathbf{A}(1-n)\bar{\mathbf{c}}+n*\mathbf{A}+\mathbf{D})(\mathbf{I}-\mathbf{A}(1-n)\bar{\mathbf{c}}+n*\mathbf{A}+\mathbf{D})^{-1})\mathbf{A}(1-n))' \\ ((\mathbf{A}(1-n)\bar{\mathbf{c}}+\mathbf{D}+n\mathbf{A})\otimes\mathbf{I}\,\mathbf{A}(1-n)\otimes\mathbf{I})' \\ \mathbf{\Sigma}_p^{-1}\mathbf{\Sigma}_{p_{t-2},p_{t-1}}^{-1}(\mathbf{A}(1-n)\mathbf{c}_{:,3}+\mathbf{D}_{:,3}+n\mathbf{A}_{:,3})'+\mathbf{\Sigma}_p^{-1}\mathbf{\Sigma}_{p_{t-2},(y_{t-1},\pi_{t-1})}(\mathbf{D}_{:,1:2}+n\mathbf{A}_{:,1:2})' \end{pmatrix} \tag{49}$$

For the mapping T to have a fixed point, all eigenvalues of $\mathbf{I}-\mathbf{A}((1-n)\mathbf{c}+n)-\mathbf{D}$ need to have positive real parts (this is equivalent to eigenvalues of $\mathbf{A}((1-n)\mathbf{c}+n)+\mathbf{D}$ being within the unit circle). Unfortunately, a full characterisation of these eigenvalues is cumbersome and provides few general insights. Importantly, the left panel of Figure 3 demonstrates that this condition is not satisfied for all values of n: the eigenvalues can turn negative if the household puts too much weight on the random walk forecasting function.

Therefore, as under inflation targeting at the zero lower bound, there exists a threshold weight $\bar{n}$. If the the household puts more weight on the random walk forecast than $\bar{n}$

[8] where $\mathbf{c}_{:,3}$ denotes the third column of $\mathbf{c}$, i.e. $\mathbf{c}_{:,3} = \tilde{\mathbf{c}}$.

Figure 3: Stability Conditions under Price Level Targeting

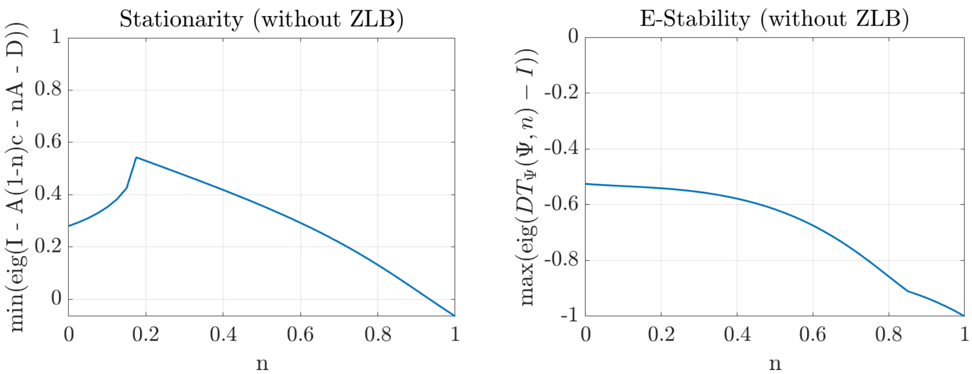

Note: This figure shows the stationary properties of the model under Price Level Targeting (i.e. the smallest real part of the eigenvalues of $I - \mathbf{A}((1-n)\mathbf{c} + n) - \mathbf{D})$ as a function of weight n put on the random walk forecasting function. For the model to be stationary all eigenvalues need to have positive real parts.

the model is not stationary anymore and the mapping T does not have a fixed point. To demonstrate how this instability property depends on the parameterization of the model, Figure 4 illustrates which combination of ϕ_π and κ lead to a stationary process $\mathbf{z}_t$ in the case of completely de-anchored expectations (i.e. $n = 1$). That is, the non-shaded area indicates the region for which the minimum eigenvalue of $I - \mathbf{A} - \mathbf{D}$ is positive. Except for the limit case of a completely flat NKPC (i.e. $\kappa = 0$), the magnitude of ϕ_π required to prevent expectations de-anchoring from causing explosive dynamics is decreasing in the slope of the Phillips curve. Importantly, expectations de-anchoring leads to explosive dynamics for empirically reasonably flat slopes of the Philips curve under Price Level Targeting unless the central bank is extremely responsive to deviations from the price level target (e.g. $\kappa = 0.2$ would require $\phi_\pi > 3.5$, $\kappa = 0.1$ would require $\phi_\pi > 6.5$).

Figure 4: Stationarity under Price Level Targeting for $n = 1$

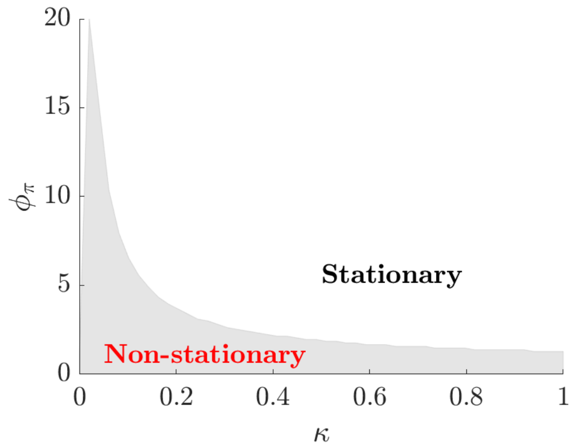

Note: This figure shows the stationarity properties of the model under Price Level Targeting under the limit case in which the household only uses the random walk forecasting function (i.e. $n = 1$) as a function of the NKPC slope κ and the central bank's responsiveness ϕ_π to price level gaps. The shaded area displays parameter combinations in which the model is non-stationary, i.e. the eigenvalues of $I - \mathbf{A} - \mathbf{D}$ are negative.

In case the model is indeed stationary and a fixed point exists, then the fixed point $\bar{\mathbf{a}}$ itself is independent of n and equivalent to the steady state under rational expectations:

$$\bar{\mathbf{a}} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{C}\bar{\mathbf{z}} = \begin{pmatrix} 0 & \frac{1-\beta}{\kappa} & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \bar{y} \\ \bar{\pi} \\ \bar{p} \end{pmatrix} \tag{50}$$

On the other hand, $\bar{\mathbf{b}}$ and $\bar{\mathbf{c}}$ are non-linear functions of n. Again, if n_t is allowed to vary over time, only $n_t = 0.5$ is consistent with a steady state in which neither heuristic makes any forecast error. In case the fixed point of $\bar{\mathbf{c}}$ exists for $n = 0.5$ (i.e. $\bar{n} \geq 0.5$), this implies an unique steady state under Price Level Targeting:

Proposition 7. *The dynamic system (41) has a unique steady state with* $\pi^* = \bar{\pi}$, $y^* = \frac{1-\beta}{\kappa}\bar{\pi}$, $r^* = \bar{\pi}$, $\hat{p}^* = 0$ *and* $n^* = 0.5$

Proof. See Appendix A.7. □

As before, for this steady state to be stable it is necessary that 1) the process $\mathbf{z}_t$ is stationary for $n = 0.5$ (i.e. $\bar{n} \geq 0.5$) so that the fixed point of the mapping T exists and that 2) the eigenvalues of

$$DT_{\mathbf{\Psi}}(\mathbf{\Psi}', n)) - \mathbf{I} \tag{51}$$

need to have negative real parts to ensure that the perceived law of motion converges to its fixed point and the steady state is learnable. As argued previously, the stationarity of the process crucially depends on the parameterization of the model so the stability of the steady state under Price Level Targeting is not a universal property. However, as the right panel of Figure 3 shows, $\bar{n} \approx 0.95$ under the estimated parameters (see next Section), so that the process $\mathbf{z}_t$ is stationary and a fixed point indeed exists for the steady state value of $n = 0.5$. In fact, the process $\mathbf{z}_t$ is stationary for all parameter combinations considered in Figure 4 when n = 0.5. Furthermore, the right panel of Figure 3 shows that the steady state is also E-stable under the estimated parameters (see next Section).

In general, under the estimated parameters, there exists a stable steady state under Price Level Targeting but the model is at risk of becoming non-stationary for high degrees of de-anchoring even away from the ZLB.

3.2.2. At the ZLB

At the ZLB the ALM in becomes

$$\boldsymbol{z}_t = \tilde{\boldsymbol{A}}\left[(1-n_t)(\boldsymbol{a} + \boldsymbol{b}\boldsymbol{w}_t + \mathbf{c}\mathbf{z}_{t-1}) + n_t \boldsymbol{z}_{t-1}\right] + \tilde{\boldsymbol{D}}\boldsymbol{z}_{t-1} + \tilde{\boldsymbol{B}}\boldsymbol{w}_t + \tilde{\boldsymbol{C}}\bar{\boldsymbol{z}} \tag{52}$$

where the matrices containing the coefficients of the ZLB ALM are denoted with a tilde. In this case no steady state exists: the price level gap at the zero lower bound follows a unit root because no deviation of inflation from target can be offset by the central bank.

Proposition 8. *Under Price Level Targeting no steady state exists at the zero lower bound*

Proof. See Appendix A.8. □

To summarise, under Price Level Targeting there is only the steady state away from the zero lower bound. Differently from the case under Inflation Targeting, this steady state is not stable for all $n \in [0, 1]$ but only for $n \in [0, \bar{n})$. The question remains whether this threshold will ever be crossed so that Price Level Targeting could lead to unstable dynamics. I will address this questions using simulation methods in Section 5.

4. Estimation and Model Dynamics

4.1. Estimation

I estimate the model using Bayesian methods to assign realistic values to the non-standard parameters θ and ω. Besides estimating the relevant parameters, this also allows me to extract an estimate of the unobserved degree of de-anchoring over time. Furthermore, Bayesian estimation provides a coherent framework to compare the data fit of the heuristic switching model with the standard rational expectations model. Due to the non-linear elements of the heuristic switching model, I compute the model likelihood $\ln p(Y \mid \Omega_{HSM})$ of observing the data Y given the parameter vector Ω_{HSM} using the bootstrap particle filter with 20,000 particles. The likelihood of the rational expectations model, on the other hand, can be computed using the Kalman filter because the model is fully linear.

The model is estimated on data from 1982Q1 to 2008Q4. I exclude the zero lower bound period because this would add yet another source of non-linearity to an already computationally complex estimation process. To match the average expectations of the heuristic switching model, I use the average 1Q ahead expectations from the Survey of Professional Forecasters for real GDP growth as well as the change of the implied GDP price deflator. To match the contemporaneous information set available to forecasters, I use the third release of the quarter-on-quarter growths rates for real GDP as well as the implied GDP price deflator. Finally, I match the model interest rate using the actual Federal Funds Rate (adjusted to quarterly frequency) since it is not subject to revisions. As in Carvalho et al. (2023), I do not use observed long-term expectations in the estimation. Instead, I will back out the model-implied long-term expectations in the next section and compare them to the available data. Therefore, the measurement equation is

$$\mathbf{Y}_t = \mathbf{\Gamma}_0 + \mathbf{\Gamma}_1 \mathbf{Z}_t + \mathbf{H}\epsilon_t \tag{53}$$

where the vector $\mathbf{Y}_t$ includes the five observable data series; $\mathbf{\Gamma}_0$ contains the means of real GDP growth, the quarterly Federal Funds Rate, the expected real GDP growth, and zeros for (expected) inflation; and the observation matrix $\mathbf{\Gamma}_1$ selects the corresponding variables from the state vector $\mathbf{Z}_t$. Since the particle filter requires some measurement error ϵ_t to avoid degeneracy, the variance $\mathbf{H}'\mathbf{H}$ of the measurement error is set to 20% of the variance of $\mathbf{Y}_t$.[9]

I use the posterior mean Ω_{RE} of the rational expectations model as a starting point to estimate the heuristic switching model in a two step procedure: I first estimate the model using the Kalman filter and take the resulting posterior as starting point for the particle filter. In each case, I use the Metropolis-Hastings algorithm with 20,000 draws and 10 separate chains to compute the full distribution.[10] [11]

[9]For comparability, I use the same assumption when estimating the RE model using the Kalman Filter.

[10]The particle filter generates random numbers at several steps of the computation so I reset the seed of the random number generator every time the likelihood is evaluated to avoid injecting randomness in the calculation of the likelihood (see e.g. Fernández-Villaverde and Rubio-Ramírez, 2007; Carvalho et al., 2023). I thereby fix the random initial conditions for the nonlinear state variables; the random draws to compute shocks in the nonlinear prediction step; and random draws in the resampling step so that the particle filter uses the same particles for every new parameter evaluation. The MCMC sampler draws from a different random number generator so that resetting the seed does not affect the parameter draws of the sampler.

[11]I scale the proposal distribution to target an overall acceptance rate of roughly 30%. Finally, I discard the first 20% of draws as burn-in, keep every 5^{th} remaining draw to compute the relevant moments, and compute the estimation results using the pooled draws of all chains.

Table 1: Prior Distributions and Posterior Estimates

Parameter		Prior			Posterior (RE)		Posterior (HSM)	
		Prior Dist.	Mean	Std.	Mean	Std.	Mean	Std.
Discount factor	β	B	0.99	0.005	0.992	0.005	0.993	0.004
IES	$\frac{1}{\sigma}$	Γ	0.5	0.2	0.155	0.071	4.771	0.600
Frisch Elast.	$\frac{1}{\eta}$	$\mathcal{N}$	0.75	0.1	0.711	0.100	1.446	0.168
Eq. Markup	$\frac{\epsilon}{\epsilon-1}$	Γ	1.35	0.2	1.326	0.084	4.887	0.445
NKPC Slope	κ	Γ	0.2	0.1	0.044	0.017	0.197	0.045
AR(1) Dem Shock	ρ_ζ	U(0,1)	0.5	-	0.675	0.061	0.510	0.119
AR(1) Mark-up shock	ρ_μ	U(0,1)	0.5	-	0.957	0.020	0.846	0.042
AR(1) MP shock	ρ_m	U(0,1)	0.5	-	0.700	0.067	0.770	0.069
Std. Demand Shock	σ_ζ	Γ^{-1}	1	0.5	0.689	0.239	1.014	0.200
Std. TFP Shock	σ_a	Γ^{-1}	1	0.5	0.877	0.308	1.084	0.198
Std. MP Shock	σ_m	Γ^{-1}	1	0.5	0.364	0.045	0.340	0.033
Feedback Output	ϕ_y	$\mathcal{N}$	0.5	0.25	0.779	0.199	0.294	0.150
Feedback Inflation	ϕ_π	$\mathcal{N}$	1.5	0.25	1.187	0.158	2.009	0.157
Learning Gain	γ	Γ	0.03	0.01	-	-	0.061	0.011
Switching Intensity	θ	Γ	5	2.5	-	-	1.882	0.357
Initial De-Anchoring	n_{1982Q1}	B	0.75	0.1	-	-	0.875	0.051
FE updating gain	ω	B	0.5	0.2	-	-	0.711	0.127
Posterior Likelihood	$\ln p(Y \mid \Omega)$	-	-	-	-88.284		-67.144	
Marginal Likelihood ($\tau = 0.5$)	$\ln p(Y)$	-	-	-	-43.078		-36.511	

Note: The marginal likelihood is estimated using Geweke's Harmonic Mean Estimator with a tuning parameter of $\tau = 0.5$. The estimates do not change significantly for different values of τ.

The independent prior distributions are described in columns 3 - 5 of Table 1. Importantly, all autoregressive parameters follow Uniform distributions centred around 0.5, with the variances following an inverse Gamma distribution centred around 1 with a variance of 0.5. For the constant gain parameter, the literature on adaptive learning suggests a range of 0.01 (see Milani, 2007) to 0.06 (see Branch and Evans, 2006). I choose an intermediate value of 0.03 as prior for γ, which follows a Beta distribution with a standard deviation of 0.01. Following the results of Cornea-Madeira et al. (2019), I impose a Gamma prior for the regime switching parameter θ that is centred at 5 with a variance of 2.5. The prior for the parameter governing the decay of forecast errors, ω, follows a relatively wide Beta distribution with mean 0.5 and variance 0.25. Finally, I estimate the initial weight of the random walk forecasting function n_{1982Q1}, assuming that inflation expectations were not perfectly anchored during the 1980s: the prior for n_{1982Q1} follows a Beta distribution centred at 0.75 with standard deviation 0.125. This initial weight in turn pins down the initial values for

$\mathbf{\Psi}_{1982Q1}$ and yields values for the initial forecast errors e^i_{1982Q1}.

The posterior distributions of the heuristic switching (rational expectations) model are presented in columns 8 & 9 (6 & 7) of Table 1. Relative to the RE model, the HSM suggests a lower real interest rate over the sample period (2.82% vs 3.23%). In line with previous literature (e.g. Ormeno, 2009; Del Negro and Eusepi, 2011; Milani, 2011), the HSM requires somewhat fewer sources of mechanical persistence than the rational expectations model, as the posterior mean of the AR(1) coefficient for demand and mark-up shocks is significantly lower (the opposite, however, is true for monetary policy shocks). On the other hand, the estimated standard deviations of demand & mar-up shocks are larger under HSM than under RE (again, this is not the case for monetary policy shocks). The posterior estimate of $\gamma = 0.061$ implies that firms and households rely on the last ~ 4 years of data. The posterior estimate of the switching intensity θ is quite a bit lower than in Cornea-Madeira et al. (2019) but, just as the memory parameter ω, remarkably close to the posterior mean of Ozden (2021). Finally, the NKPC slope estimate under rational expectations is much larger than under forecast switching.

At the posterior mean, the HSM model fits the observed data better than the RE model. The same still holds for the overall model fit (i.e. accounting for the additional degrees of freedom stemming from four extra parameters by marginalising out their influence). The estimates for marginal likelihood $\ln p(Y \mid \mathcal{M})$ in the lower panel of Table 1 suggest that the heuristic switching model fits the data much better than the standard rational expectations model, which is a common feature of models with non-rational expectations (Milani, 2007).[12]

4.2. Model Predictions

I back out the model-implied states using the nonlinear smoother of Godsill et al. (2004). To document the recent performance of the model, I also compute the out-of-sample model forecasts up until 2023Q2. That is, I estimate the model on the sample from 1982Q1:2008Q4 and use the resulting parameter estimates with an extended set of observables that includes shadow interest rate of Wu and Xia (2016) to compute the out-of-sample model predictions until 2023Q2.[13] The resulting smoothed state estimates are plotted (in blue) along with the data (in black) in Figure 5. To account for the 2020Q1-2020Q4 data outliers driven by the Covid-19 pandemic and the resulting policy measures, I set the measurement error in these 4 quarters equal to the variance of the observed data over the same period.

[12] I compute the marginal Likelihood using the Modified Harmonic Mean Estimator of Geweke (1999).

[13] The observable data after 2008Q4 is demeaned using a separate mean different from the estimation sample, implying a structural break. Furthermore, the measurement error after 2008Q4 is adjusted to reflect the changed data variance during that period.

Figure 5: Model Predictions

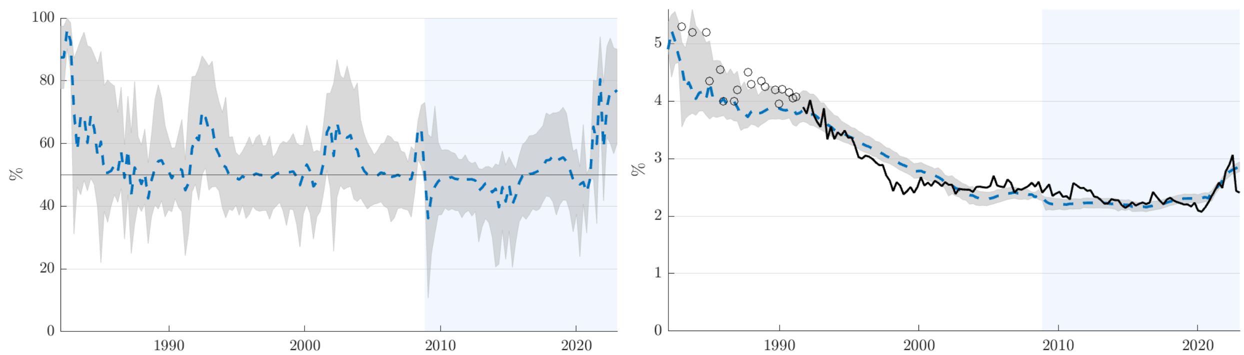

(a) n_t (b) 10-year Inflation Expectations

— SPF ○ Blue Chip & Livingston Survey - - - HSM (with 95% CI)

Note: The panels displays the smoothed model states. The solid black line denotes the data; the blue dashed line shows the median model prediction of the heuristic switching model along with 95% confidence intervals in grey. The blue-shaded area denotes the out-of-sample forecast period.

As the left Panel of Figure 5 shows, the implied weight of the random walk forecasting function – the gauge of expectations de-anchoring – is quite elevated in the early 1980s before the Volcker driven re-anchoring was completed. Since the late 1980s the weight of the random walk forecasting function has largely fluctuating around its equilibrium value of 50%, with spikes to the upward coinciding with the recessionary periods of the early 1990s and 2000s. Interestingly, and contrary to the findings of Gáti (2023), the model does not detect any significant de-anchoring following the global financial crisis. In recent years, the model implies significant expectations de-anchoring beginning in 2021 as the result of large, positive inflation forecast errors.

Furthermore, despite not being used as observable data in the estimation, the model-implied long-run inflation expectations match the SPF forecasts quite well (see the right panel of Figure 5). They follow the steady downward trend in the first half of the sample and capture the relatively fast increase in the early 2020s. As a consequence, the sensitivity of model-implied inflation forecasts to forecast errors is quite close to the sensitivity present in the SPF forecasts (see Figure B.2). This illustrates the implicit mapping to the endogenous gain approach of Carvalho et al. (2023), as the aggregate expectations under the heuristic switching model exhibit a time-varying gain in the aggregate. Finally, the model-implied IRF of long-term inflation expectations to a 1pp monetary policy shock matches the IRF estimated on observed expectations rather well (see Figure C.1).

In general, neither the rational expectations model nor the heuristic switching model captures all the high-frequency fluctuations of the observable data as they are attributed to measurement error. Instead, both models track the apparent underlying low-frequency movements quite well (see Figure B.1). Importantly, the (time-varying) measurement error almost completely smoothes out the huge spikes in GDP growth during the Covid-19 pandemic.

4.3. The Effects of De-anchoring

The weight that the representative household puts on the random walk forecasting function has a significant impact on the dynamics of the model. To illustrate this, Figure C.3 plots the impulse response functions of output, inflation, and interest rates to a demand, supply, and interest rate shock, respectively, for varying weights.[14] The responses to all three shocks clearly become more persistent the higher the degree of de-anchoring is. Furthermore, the IRFs display an increasingly hump shaped pattern for higher degrees of de-anchoring, culminating in an oscillating response for fully naive expectations. This is part of a broader pattern: the variance of output and inflation of the heuristic switching model is weakly higher than under rational expectations for all degrees of de-anchoring (and increasing therein, see Figure C.4). Only the covariance between output and inflation can turn below its rational expectations counterpart for very high degrees of de-anchoring as both variables increasingly approach a random walk on their own.

Even more drastic, however, is the consequence of de-anchoring when the zero lower bound is potentially binding (see Figure 6): A sequence of four consecutive σ_ζ discount factor shocks decreases demand enough to force interest rates to the zero lower bound and leads to significant drops in output and inflation. Upon impact of the sqeuence of contractionary shocks, the adaptive learning rule initially predicts the resulting contraction better than the random walk forecast as the latter adjusts expectations only with a lag. As the shocks start to level off, however, the adaptive forecasting function yields worse forecasts. This is the case because, coming out of steady state with an active monetary policy response, it fails account for the fact that the central bank is constrained in its ability to stabilise the economy. That is, because the adaptive forecasting function starts with $\bar{\boldsymbol{\Psi}}_0$ at its steady state value where where monetary policy can mediate the effect of exogenous shocks, it fails to accurately reflect the new environment the ZLB is binding. Therefore, the household (rightly) starts to believe that the central bank is not able to stabilise inflation and thus adopts a random walk forecasting function, reflecting this view. This ultimately pushes the economy into a deflationary spiral.

Note that, for the first 4 periods, the HSM model yields a qualitatively similar responses

[14] That weight is held fixed for the horizon of the IRF

Figure 6: De-anchoring at the ZLB

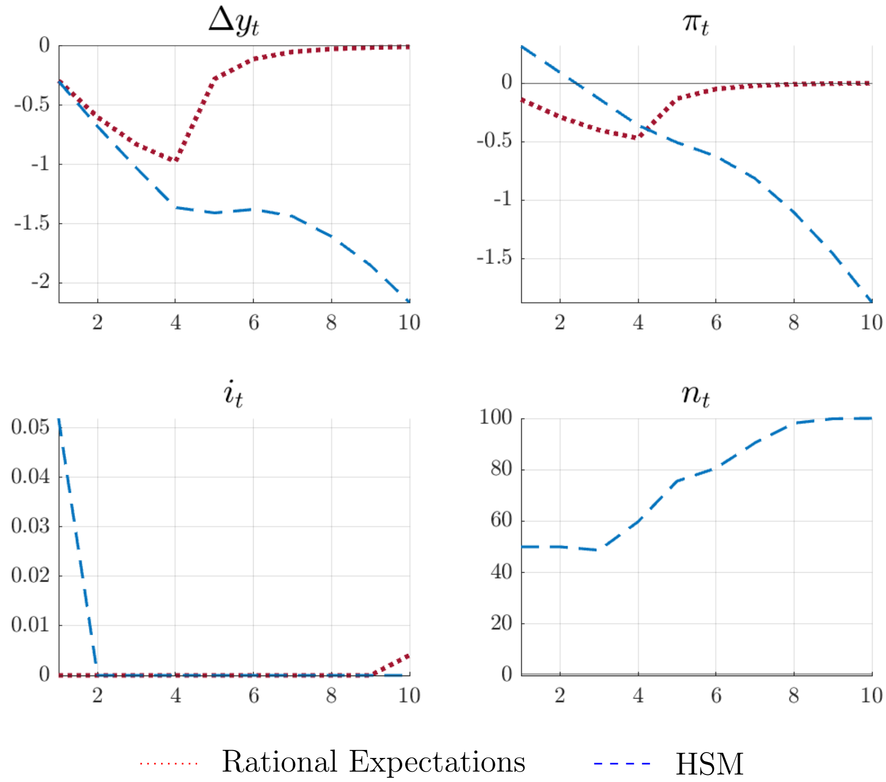

Note: IRF to four consecutive standard deviation shocks to the discount factor ζ_t. The dashed blue line shows the IRFs of the heuristic switching model (HSM); the dotted red line shows the IRFs of the rational expectations model (RE).

as the rational expectations model (plotted in red).[15] The key difference is that forecast switching can generate endogenous, belief-driven deflationary spirals at the zero lower bound. The rational expectations model, on the other hand, returns to the steady state once the initial shocks dissipate and the central bank leaves the zero lower bound after 9 quarters.

The deflationary spiral is driven by the forecast switching behaviour. In fact, without forecast switching (i.e. keeping $n = 0.5$ constant), the economy would very slowly return to its pre-crisis level as the counterfactual analysis in Figure C.6 shows. Note that this scenario occurs only when the central bank is constrained by the zero lower bound. If monetary policy is unconstrained, the system remains stable even when the households relies only on the representative forecasting function (i.e. $n = 1$). However, even when monetary policy remains unconstrained and can prevent deflationary spirals, the economy still experiences elevated volatility of output and inflation. Therefore, the potential downside posed by de-anchoring is highly asymmetric with regard to the proximity to the ZLB, i.e. depends on

[15]The rational expectations IRF is computed using the parameters of the HSM posterior for better comparability.

the level of the equilibrium interest rate.

5. Policy Implications

As discussed in the previous section, sequences of adverse shocks can push the economy to the zero lower bound and lead to deflationary spirals by causing expectations to de-anchor. How can monetary policy prevent this scenario? In this section, I first find the optimal Taylor rule coefficient on inflation numerically. I then compare the optimal inflation targeting Taylor rule with a Price Level Targeting rule, extending the analytical discussion of Section 3.2.

5.1. The Optimal Taylor Rule

To find the optimal Taylor rule coefficient on inflation ϕ_π^*, I simulate the model over 2,000 periods over a grid of $\Phi_\pi \in [1, 50]$, holding all other parameters fixed at the posterior mean of the heuristic switching model. The optimal Taylor rule coefficient is the grid point that maximises the household's utility (see Equation (1)) in a cross-section of 200 simulations. Importantly, I simulate the non-linear model (see Section 2.1) instead of the linearised model (see Section 3.1.1)[16] to capture the non-linearities arising from the ZLB and from the potential of deflationary spirals.[17] I consider two different equilibrium interest rate cases: First, I set the discount factor β such that the equilibrium nominal interest rate is $\sim 6\%$ (assuming and inflation target of 2%), as was prevailing in the early 1980s according to the New York Fed's Holston-Laubach-Williams model. Second, I increase the discount rate such that the equilibrium nominal interest rate decreases to $\sim 3\%$. Furthermore, in light of the results of Carvalho et al. (2023) and Gáti (2023), I also consider a learning setup (denoted by $\widehat{\text{HSM}}$) in which the forecast error of interest rate expectations are not used to evaluate the two forecasting functions.[18] This alternative setup allows the central bank to respond much more aggressively to inflation because forecast errors on interest rates do not lead to further de-anchoring by themselves.

Table 2 presents the optimal inflation response in the six different cases. In a high

[16]The non-linear model is solved by finding the labour and inflation policy functions that set the error of the two inter-temporal conditions (3) & (10) to zero given the predetermined subjective expectations. For tractability I maintain the assumption that households use linear predictors to forecast the *deviations* of output and inflation from steady state but translate these expectations back into *levels* of output and inflation to solve the non-linear model. I initialise the simulation at the fixed point of the adaptive forecasting function. In the case of Price Level Targeting, the fixed point is derived under the assumption of no autocorrelation in the shock processes. Therefore, I simulate the linearised model first to obtain the starting values for the nonlinear simulation.

[17]For simplicity, I impose market clearing to derive the consumption forecast for the Euler equation (3).

[18]That is, Equation (21) only uses the squared forecast errors for inflation and output

Table 2: Optimal Inflation Response - Simulation Evidence

	R*=6%			R*=3%		
	RE	HSM	$\widehat{\text{HSM}}$	RE	HSM	$\widehat{\text{HSM}}$
ϕ_π^*	23.94	17.53	55.39	20.33	27.34	41.87

Note: This table shows the optimal inflation response in a standard Taylor rule based on a simulation of the non-linear model in high nominal interest rate environment (i.e. $R^* \sim 6\%$) and under a low nominal interest rate environment (i.e. $R^* \sim 3\%$) for the Rational Expectations (RE), model the Heuristic Switching Model (HSM), and the Heuristic Switching Model without the evaluation of Interest Rate expectations ($\widehat{\text{HSM}}$).

interest rate environment, the optimal response to inflation is lower under the heuristic switching model than under rational expectations. This is the case because the forecast errors associated with aggressive monetary policy moves can lead to expectations de-anchoring, which in turn leads to a higher volatility of output and inflation. When this channel is shut down, the optimal interest rate response under the heuristic switching model becomes much larger than under rational expectations, mirroring the results of Carvalho et al. (2023) and Gáti (2023). That is, a central bank faces the trade-off between missing its inflation target, which can give rise to expectations de-anchoring, and extreme interest rate adjustments which, depending on the expectation formation, can equally give rise to expectations de-anchoring. This non-linear relationship between the central bank's responsiveness to inflation and the degree of de-anchoring is illustrated in Figure 7. Conversely, in a lower interest rate environment, the optimal inflation response is much higher than under rational expectations regardless of the assumption regarding interest rate forecast errors. That is, a central bank is willing to accept some degree of de-anchoring caused by extreme interest rate adjustments to prevent the zero lower bound. This result illustrates the magnitude of the welfare loss of de-anchoring at the zero lower bound and shows the importance of accounting for the ZLB when discussing the consequences of de-anchoring.

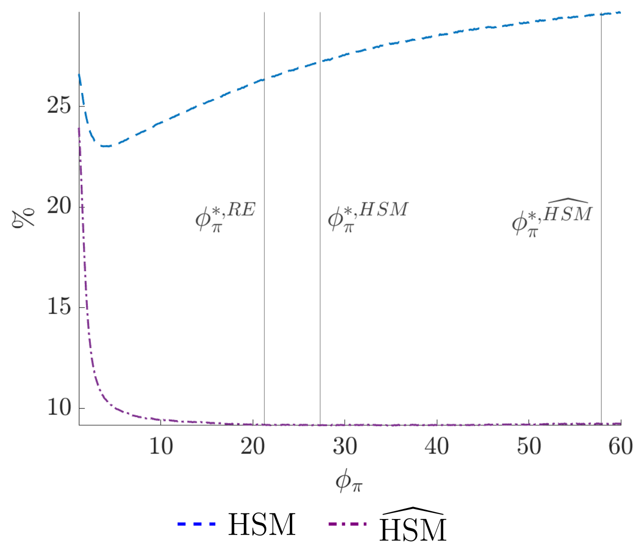

Note: This figure displays the variance of weight n_t put on the random walk forecasting function for different degrees of responsiveness to inflation under an Inflation Targeting framework away from the zero lower bound with and without evaluation of interest rate forecast errors (dashed and dash-dotted lines, respectively). The optimal inflation response under Rational Expectations, the Heuristic Switching Model, and the Heuristic Switching Model without the evaluation of Interest Rate expectations is denoted by ϕ_π^{RE}, ϕ_π^{HSM}, $\phi_\pi^{\widehat{HSM}}$, respectively.

Figure 7: Weight of Random Walk Forecast - Variance

5.2. *Other Monetary Policy Frameworks*

To supplement the analytical discussion of Price Level Targeting and to evaluate its performance vis-à-vis Inflation Targeting, I simulate the model economy under the two different policy regimes in the same two interest rate environments as above.

Table 3: Policy Framework Evaluation - Simulation Evidence

	$R^* = 6\%$		$R^* = 3\%$	
	IT	PLT	IT	PLT
Welfare $\mathcal{W}$	-126.48	-125.97	-385.77	-416.88
	[-192.42 -106.73]	[-197.41 -105.46]	[-566.56 -326.38]	[-113,988.11 -335.19]
Risk of Defl. Spiral (%)	0.13	0.25	0.50	0.75
	[0.00 0.47]	[0.13 0.75]	[0.13 1.38]	[0.13 1.50]

Note: This table shows the simulation results under inflation targeting (IT) and price level targeting (PLT) in a high nominal interest rate environment (i.e. $R^* \sim 6\%$) and a low nominal interest rate environment (i.e. $R^* \sim 3\%$). Values in square brackets indicate 95% confidence bands, which are computed by taking 500 parameter draws from the draws of the MCMC sampler.

Table 3 presents the simulation results. In a high nominal interest rate world, Price

Level Targeting leads to superior outcomes in terms of welfare relative to Inflation targeting. This result switches in a world of low nominal interest rates. For both policy frameworks it becomes significantly harder to stabilise the economy due to a potentially binding zero lower bound. The constraint on interest rate setting leads to significantly more variation in the degree of de-anchoring. These two factors, i.e. the zero lower bound and the associated higher risk of de-anchoring, increase the volatility of output and inflation. The resulting drop in welfare is much more pronounced under Price Level Targeting. To understand why this is the case, consider a deflationary shock that pushes the economy to the ZLB. This shock crates large forecast errors because it cannot be offset by the central bank under either policy framework. These forecast errors are even larger under Price Level Targeting because such a contractionary shock would normally be associated with an increase in inflation expectations generated by PLT to compensate for the price level gap. However, due to the ZLB, this overshooting never materialises. Therefore, the agent increasingly believes to be in a world with a passive central bank and accordingly adopts the random-walk forecasting function. Finally, de-anchoring under Price Level Targeting increases the variance of output and inflation much more than under Inflation targeting as Figure C.5 illustrates. Therefore, the main driver of de-anchoring under Price Level Targeting is the failure to generate offsetting inflation which causes a rapid shift to a random walk forecasting function. This shift is much more detrimental under Price Level Targeting than under Price Level Targeting.

These results hold under a wide range of parameter combinations. In fact, in a subsample of 500 draws from the MCMC sampler, 81% of the parameter combinations imply that Inflation Targeting outperforms Price Level Targeting in a low interest rate environment (compared to 16% in a high interest rate environment). This is illustrated by the very negative left tail of the distribution of simulated welfare outcomes shown in the square brackets of Table 3.

The inability of Price Level Targeting to address the threat of expectations de-anchoring, differs from the main result of Honkapohja and Mitra (2020) for two main reasons. First, Honkapohja and Mitra (2020) calibrate a relatively small constant gain parameter of 0.005. The gain under imperfect credibility, therefore, is much smaller than in my case, where the random walk forecast of course implies a gain of one. The second main reason is that their model does not feature supply shocks. The latter reason is part of a broader point: Models of de-anchoring developed in a time of low nominal interest rate and relatively subdued supply shocks (such as Honkapohja and Mitra (2020), Bianchi et al. (2021), and an earlier version of this paper) generally favour alternatives to inflation targeting. This recommendation changes, however, once supply shocks are taken into account, the proposed policy alternatives can actually increase the risk of de-anchoring instead of reducing it (see

Section E for a discussion of the asymmetric Taylor rule proposed by Bianchi et al. (2021) in the face of endogenously anchored expectations).

In a final step of this simulation exercise, I confirm the optimal policy recommendation of Gáti (2023): a central bank can improve welfare by reacting forcefully to movements in long-term inflation expectations. This is shown by the welfare comparison between optimal Inflation Targeting as well as an optimal Taylor rule with long-run inflation expectations as an additional input in Table D.1. Responding to deviations of long-run inflation expectations from the inflation target allows the central bank to accommodate inflation fluctuations when expectations are well-anchored, while implicitly conditioning its reaction function on the degree of de-anchoring. Interestingly, this result arises in a model of 1 period ahead expectations and does not require the infinite horizon formulation employed by Gáti (2023). However, my simulated model suggests that the central bank needs to be much less aggressive to achieve this goal than suggested by Gáti (2023).

6. Conclusion

Central bankers frequently voice concerns about the possibility of de-anchored inflation expectations, that is, the risk that short-term developments in inflation feed into long-term inflation expectations. The contribution of this paper is to study the anchoring of inflation expectations. For this purpose, I build a model in which a representative household weighs his beliefs about whether or not the Taylor principle is satisfied. Given this belief, he either forms expectations using a linear adaptive learning function or, when he doubts the central bank's commitment, using a random walk forecasting function. The time-varying weight put on the naive forecasting heuristic determines the sensitivity of short- and long-run expectations to short-run conditions. The model has the same steady states as under rational expectations but features complex dynamics away from the steady state that are non-linear in the degree of anchoring: when expectations de-anchor, the volatility of output growth and inflation increases. Monetary policy can prevent expectations de-anchoring from causing inflationary or deflationary spirals if the Taylor principle is satisfied and the zero lower bound is not binding. At the zero lower bound, however, de-anchoring can lead to a self-fulfilling deflationary spiral. Thus, the potential welfare loss of de-anchoring is asymmetric and bigger in a low interest rate environment. I estimate the model using the non-linear particle filter on U.S. data and use the estimated model to explore the implications for monetary policy, in both a high and a low nominal interest rate environment. A striking result is that price level targeting can be de-stabilising when employed near the zero lower bound, due to its more restrictive stability requirements. However, in a high nominal interest rate environment, it successfully stabilises the economy.

References

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2020). Imperfect Macroeconomic Expectations: Evidence and Theory. In *NBER Macroeconomics Annual 2020, volume 35*, pp. 1–86. University of Chicago Press.

Anufriev, M., T. Assenza, C. Hommes, and D. Massaro (2013). Interest rate rules and macroeconomic stability under heterogeneous expectations. *Macroeconomic Dynamics 17*(8), 1574–1604.

Beechey, M. J., B. K. Johannsen, and A. T. Levin (2011, April). Are long-run inflation expectations anchored more firmly in the euro area than in the united states? *American Economic Journal: Macroeconomics 3*(2), 104–29.

Benhabib, J., S. Schmitt-Grohe, and M. Uribe (2002). Avoiding Liquidity Traps Avoiding Liquidity Traps. *Journal of Political Economy 110*(3), 535–563.

Bianchi, F., L. Melosi, and M. Rottner (2021). Hitting the Elusive Inflation Target. *Journal of Monetary Economics 124*, 107–122.

Branch, W. A. and G. W. Evans (2006). A simple recursive forecasting model. *Economics Letters 91*(2), 158–166.

Brock, W. A. and C. H. Hommes (1997). A Rational Route to Randomness. *Econometrica 65*(5), 1059–1096.

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

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023). Anchored inflation expectations. *American Economic Journal: Macroeconomics 15*(1), 1–47.

Cornea-Madeira, A., C. Hommes, and D. Massaro (2019). Behavioral Heterogeneity in U.S. Inflation Dynamics. *Journal of Business and Economic Statistics 37*(2), 288–300.

De Grauwe, P. and Y. Ji (2019). Inflation Targets and the Zero Lower Bound in a Behavioural Macroeconomic Model. *Economica 86*(342), 262–299.

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

Eusepi, S., M. Giannoni, and B. Preston (2020a). On the Limits of Monetary Policy.

Eusepi, S., M. Giannoni, and B. Preston (2020b). On the limits of monetary policy. In *NBP Summer Workshop Conference paper*.

Evans, C. L. (2012). Monetary Policy in a Low-Inflation Environment: Developing a State-Contingent Price-Level Target. *Journal of Money, Credit and Banking 44*(s1), 147–155.

Evans, G. W. and S. Honkapohja (2001). *Learning and expectations in macroeconomics*. Princeton University Press.

Fernández-Villaverde, J. and J. F. Rubio-Ramírez (2007). Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies 74*(4), 1059–1087.

Gáti, L. (2023). Monetary policy & anchored expectations—an endogenous gain learning model. *Journal of Monetary Economics*.

Geweke, J. (1999). Using simulation methods for bayesian econometric models: Inference, development, and communication. *Econometric Reviews 18*(1), 1–73.

Godsill, S. J., A. Doucet, and M. West (2004). Monte carlo smoothing for nonlinear time series. *Journal of the American Statistical Association 99*(465), 156–168.

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

Gürkaynak, R. S., B. Sack, and E. Swanson (2005). The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American Economic Review 95*(1), 425–436.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2022, 02). The Slope of the Phillips Curve: Evidence from U.S. States*. *The Quarterly Journal of Economics 137*(3), 1299–1344.

Honkapohja, S. and K. Mitra (2020). Price level targeting with evolving credibility. *Journal of Monetary Economics 116*(C), 88–103.

Kumar, S., H. Afrouzi, O. Coibion, and Y. Gorodnichenko (2015). Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand. NBER Working Paper 21814, National Bureau of Economic Research.

Lansing, K. J. (2021). Endogenous Forecast Switching Near the Zero Lower Bound. *Journal of Monetary Economics 117*, 153–169.

Levin, A. T., F. M. Natalucci, and J. M. Piger (2004). The Macroeconomic Effects of Inflation Targeting. *Review - Federal Reserve Bank of Saint Louis 86*(Jul), 51–80.

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

Marcet, A. and J. P. Nicolini (2003). Recurrent hyperinflations and learning. *American Economic Review 93*(5), 1476–1498.

Mele, A., K. Molnár, and S. Santoro (2020). On the perils of stabilizing prices when agents are learning. *Journal of Monetary Economics 115*(C), 339–353.

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

Milani, F. (2011). Expectation shocks and learning as drivers of the business cycle. *The Economic Journal 121*(552), 379–401.

Ormeno, A. (2009). Disciplining expectations: adding survey expectations in learning models. 2009 Meeting Papers 1140, Society for Economic Dynamics.

Ozden, T. (2021). Heterogeneous Expectations and the Business Cycle at the Effective Lower Bound. DNB Working Papers 714, De Nederlandsche Bank.

Ozden, T. and R. Wouters (2021). Restricted Perceptions, Regime Switches and the Effective Lower Bound.

Reis, R. (2021). Losing the Inflation Anchor. CEPR Discussion Papers 16664, Centre for Economic Policy Research.

Romer, C. D. and D. H. Romer (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review 94*(4), 1055–1084.

Rotemberg, J. J. (1982). Sticky Prices in the United States. *Journal of Political Economy 90*(6), 1187–1211.

Strohsal, T., R. Melnick, and D. Nautz (2016). The time-varying degree of inflation expectations anchoring. *Journal of Macroeconomics 48*(C), 62–71.

Wu, J. C. and F. D. Xia (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking 48*(2-3), 253–291.

Appendix A. Derivation of the Fixed Point

Take the ODE

$$\frac{\partial \mathbf{\Psi}'}{\partial \tau} = \mathbf{R}^{-1} E\mathbf{x}_{t-2}\left(\mathbf{A}\left[(1-n)\mathbf{\Psi}\mathbf{x}_{t-1})\right] + \mathbf{A}n\mathbf{z}_{t-2} + \mathbf{B}\mathbf{w}_t + \mathbf{C}\bar{\mathbf{z}}) - \mathbf{\Psi}\mathbf{x}_{t-2}\right)' \tag{54}$$

and work separately with $\mathbf{a}, \mathbf{b}$. Given that $\lim_{t\to\infty} E\mathbf{x}_t\mathbf{x}_t\prime = \Sigma_x = \mathbf{R} = diag(1, \mathbf{\Sigma_w})$, it is possible to write the ODE as

$$\frac{\partial \mathbf{a}'}{\partial \tau} = (\mathbf{A}\left[(1-n)\mathbf{a}\right] + n\mathbf{A}\mathbf{z}_{t-2} + \mathbf{C}\bar{\mathbf{z}})' - \mathbf{a}' \tag{55}$$

$$\frac{\partial \mathbf{b}'}{\partial \tau} = \mathbf{\Sigma_w}^{-1}\mathbf{\Sigma}_{\mathbf{w}_{t-2},\mathbf{w}_{t-1}}(\mathbf{A}(1-n)\mathbf{b} + \mathbf{B})' + \mathbf{\Sigma}_w^{-1}\mathbf{\Sigma}_{\mathbf{w}_{t-2},\mathbf{z}_{t-2}}(\mathbf{A}n)' - \mathbf{b}' \tag{56}$$

As long as the ALM in Equation (28) is asymptotically stationary, i.e.

$$\lambda < 1 \quad \forall \lambda \in \Lambda = \{\Lambda : |\mathbf{I} - n\mathbf{A} - \mathbf{\Lambda}\mathbf{I}| = 0\} \tag{57}$$

we can write, holding $(\mathbf{\Psi}', n)$ fixed, $\lim_{t\to\infty} E\mathbf{z_t}$ as

$$\mathbf{z}_t(\mathbf{\Psi}, n) = \left((\mathbf{I} - n\mathbf{A})^{-1}\left(\mathbf{A}\left[(1-n)\mathbf{a}\right] + \mathbf{C}\bar{\mathbf{z}}\right)\right)' \tag{58}$$

so that the ODE for $\mathbf{a}$ becomes

$$\frac{\partial \mathbf{a}'}{\partial \tau} = \left(\mathbf{A}\left[(1-n)\mathbf{a} + n(\mathbf{I} - n\mathbf{A})^{-1}\left(\mathbf{A}(1-n)\mathbf{a} + \mathbf{C}\bar{\mathbf{z}}\right)\right] + \mathbf{C}\bar{z}\right)' - \mathbf{a}' \tag{59}$$

yielding the fixed point

$$vec(\bar{\boldsymbol{a}}') = (\boldsymbol{I} - \boldsymbol{A})^{-1} vec\left((\boldsymbol{C}\bar{\boldsymbol{z}})'\right) \tag{60}$$

Turning our attention to the ODE for $\mathbf{b}$, we note that $\mathbf{\Sigma}_{\mathbf{w}_{t-2},\mathbf{w}_{t-1}} = \mathbf{F}\mathbf{\Sigma_w} = \mathbf{\Sigma_w}\mathbf{F}$ (because both are diagonal) and that $\Sigma_{\mathbf{w}_{t-2},\mathbf{z}_{t-2}} = \Sigma_{\mathbf{w}_t,\mathbf{z}_t}$ is endogenously determined:

$$\mathbf{\Sigma}_{\mathbf{w}_t,\mathbf{z}_t} = E(\mathbf{w}_t)(\mathbf{A}\left[(1-n)\mathbf{\Psi}\mathbf{x}_t)\right] + \mathbf{A}n\mathbf{z}_{t-1} + \mathbf{B}\mathbf{w}_{t-1} + \mathbf{C}\bar{z})' \tag{61}$$

$$\Leftrightarrow \mathbf{\Sigma}_{\mathbf{w}_t,\mathbf{z}_t} = \Sigma_\mathbf{w}\mathbf{b}'(\mathbf{A}(1-n))' + \mathbf{\Sigma_w}\mathbf{B}' + \Sigma_{\mathbf{w}_t,\mathbf{z}_{t-1}}(\mathbf{A}n)' \tag{62}$$

where $\Sigma_{\mathbf{w}_t,\mathbf{z}_{t-1}}$ again is endogenous

$$\mathbf{\Sigma}_{\mathbf{w}_t,\mathbf{z}_{t-1}} = E(\mathbf{F}\mathbf{w}_{t-1})(\mathbf{z}_{t-1})' = \mathbf{F}\mathbf{\Sigma}_{\mathbf{w_t},\mathbf{z_t}} \tag{63}$$

so that

$$vec(\mathbf{\Sigma}_{\mathbf{w}_t,\mathbf{z}_t}) = (\mathbf{I} - \mathbf{A}n \otimes \mathbf{F})^{-1}\left(\mathbf{A}(1-n) \otimes \mathbf{\Sigma_w} vec(\mathbf{b}') + vec(\mathbf{\Sigma_w}\mathbf{B}')\right) \tag{64}$$

Substituting this into into $vec(^{\partial \mathbf{b}'}/_{\partial \tau})$

$$vec\left(\frac{\partial \mathbf{b}'}{\partial \tau}\right) = \mathbf{A}n \otimes \mathbf{\Sigma}_w^{-1}(\mathbf{I} - \mathbf{A}n \otimes \mathbf{F})^{-1}\left(\mathbf{A}(1-n) \otimes \mathbf{\Sigma_w} vec(\mathbf{b}') + vec(\mathbf{\Sigma_w}\mathbf{B}')\right) \tag{65}$$

$$+ \mathbf{A}(1-n) \otimes \mathbf{F} vec(\mathbf{b}') + vec(\mathbf{F}\mathbf{B}') - vec(\mathbf{b}') \tag{66}$$

which, as $\tau \to 0$, yields the fixed point

$$vec(\bar{\boldsymbol{b}}') = (\boldsymbol{I} - \mathbf{G}_1)^{-1}\, \mathbf{G}_2 vec(\mathbf{B}') \tag{67}$$

where

$$\mathbf{G}_1 = \mathbf{A}(1-n) \otimes \mathbf{F} + \mathbf{A}n \otimes \mathbf{\Sigma_w}^{-1} (\boldsymbol{I} - n\mathbf{A} \otimes \mathbf{F})^{-1} \mathbf{A}(1-n) \otimes \mathbf{\Sigma_w} \tag{68}$$

$$\mathbf{G}_2 = \mathbf{I} \otimes \mathbf{F} + \mathbf{A}n \otimes \mathbf{\Sigma_w}^{-1} (\boldsymbol{I} - n\mathbf{A} \otimes \mathbf{F})^{-1} \mathbf{I} \otimes \mathbf{\Sigma_w} \tag{69}$$

The mapping T from the perceived to the actual law of motion for the PLM coefficients for a given n can therefore be characterised as

$$T(\mathbf{\Psi}', n) = \begin{pmatrix} T(\mathbf{a}', n) \\ T(\mathbf{b}', n) \end{pmatrix} = \begin{pmatrix} ((\boldsymbol{I} - \boldsymbol{A}n)^{-1}\boldsymbol{A}(1-n))' \\ \mathbf{G}_1 \end{pmatrix} \tag{70}$$

A.1. Proof of Proposition I

Proof. A fixed point exists if the process z_t is asymptotically stationary. Since the eigenvalues of $(n\mathbf{A})$ are increasing in n, a fixed point exists if the eigenvalues of A lie within the unit circle. Bullard and Mitra (2002) show that a necessary and sufficient condition for the eigenvalues of A to lie within the unit circle for a bivariate Taylor rule is

$$\kappa\,(\phi_\pi - 1) + (1-\beta)\phi_y > 0 \tag{71}$$

Strict inflation targeting is the limit case with $\phi_y = 0$, thus reducing the requirement to

$$\phi_\pi > 1 \tag{72}$$

i.e. the standard Taylor principle. □

A.2. Proof of Proposition II

Proof. I guess that the fixed point of the mapping T is a steady state, , i.e. $\mathbf{z}^* = \bar{\mathbf{a}} = (\frac{1-\beta}{\kappa}\bar{\pi}, \bar{\pi})'$. In the absence of structural shocks and given $\mathbf{z}_t = \mathbf{z}_{t-1}$, neither forecast heuristic makes prediction errors so that $n_t = n = 0.5$ and the system becomes

$$y^* = (1-n)y^* + ny^* - \frac{1}{\sigma}\left(\bar{\pi} + \phi_\pi(\pi^* - \bar{\pi}) - (1-n)\pi^* - n\pi^*\right) \tag{73}$$

$$\pi^* = \beta[(1-n)\bar{\pi} + n\pi^*] + \kappa y^* \tag{74}$$

It is clear that $\pi^* = \bar{\pi}$ and $y^* = \frac{1-\beta}{\kappa}\bar{\pi}$ is a solution, thus confirming the initial guess. This is the case if the weight of the random walk forecast is fixed at any value $n_t = n \in (0,1)$. However, due to the decay of past prediction errors and the forecast selection mechanism, only $n_t = n^* = 0.5$ can be a steady state. □

A.3. *Proof of Proposition III*

Proof. When all households are adaptive learners (i.e. $n_t = 0$), $DT_a(\mathbf{a}', n)$ reduces to $[\mathbf{A}]$. Given that the eigenvalues of $\mathbf{A}$ are smaller than one in absolute value, the eigenvalues of $[\mathbf{A} - \mathbf{I}]$ will be negative since we subtract 1 from each of the eigenvalues of $\mathbf{A}$. Similarly, $DT_b(\mathbf{b}', n)$ reduces to $[\mathbf{A} \otimes \mathbf{F}]$, so that the eigenvalues of $[\mathbf{A} \otimes \mathbf{F} - \mathbf{I}]$ will be negative as well. When all households are naive forecasters (i.e. $n_t = 1$), both $DT_a(\mathbf{a}', n)$ and $DT_b(\mathbf{b}', n)$ reduce to zero so that the associated eigenvalues are of course -1. Any combination in between for $n \in (0, 1)$ is a non-linear combination of the two preceding cases whose eigenvalues are decreasing in n and therefore still features negative eigenvalues, as Figure 2 illustrates. □

A.4. *Proof of Proposition IV*

Proof. The eigenvalues of $\tilde{\mathbf{A}}$ are given by

$$\lambda_{1,2} = \frac{1}{2}\left(1 + \beta + \frac{\kappa}{\sigma} \pm \sqrt{\left(1 + \beta + \frac{\kappa}{\sigma}\right)^2 - 4\beta}\right) \tag{75}$$

which implies for all positive values of β, κ, and σ that $|\lambda_1| < 1$ and $\lambda_2 > 1$. Since the eigenvalues of $n\tilde{\mathbf{A}}$ are increasing in n, there exists $\bar{n} = \frac{2}{\left(1+\beta+\frac{\kappa}{\sigma}+\sqrt{(1+\beta+\frac{\kappa}{\sigma})^2-4\beta}\right)} \in (0, 1)$ such that only for $n < \bar{n}$ all eigenvalues of $n\tilde{\mathbf{A}}$ fall within the unit circle, the process $\mathbf{z}_t$ is asymptotically stationary, and the mapping $\tilde{T}$ has a fixed point. For $n \geq \bar{n}$ the process is not stationary and no fixed point exists. □

A.5. *Proof of Proposition V*

Proof. I again guess that $\tilde{\mathbf{z}}^* = \tilde{\bar{\mathbf{a}}} = (0, 0)'$. In the absence of structural shocks and given $\mathbf{z}_t = \mathbf{z}_{t-1}$, neither forecast heuristic makes prediction errors so that $n_t = n = 0.5$ and the system becomes

$$\tilde{y} = (1 - n)\tilde{y} + n_t\tilde{y} - \frac{1}{\sigma}\left(-((1 - n)\tilde{\pi} + n\tilde{\pi})\right) \tag{76}$$

$$\tilde{\pi} = \beta((1 - n)\tilde{\pi} + n\tilde{\pi}) + \kappa\tilde{y} \tag{77}$$

which is solved by $\tilde{\pi}^* = \tilde{y}^* = 0$, confirming the initial guess. □

A.6. *Proof of Proposition VI*

Proof. First, I assume that $\bar{n} < 0.5$. In this case the steady state with $n^* = 0.5$ cannot be stable because the process $\mathbf{z}_t$ is non-stationary and $\tilde{\mathbf{a}}_t$ will never converge to any fixed point

$\tilde{\bar{\mathbf{a}}} = \mathbf{z}^*$.

Second, I assume that $\bar{n} > 0.5$ and consider the two extreme cases of n_t: when all households are adaptive learners (i.e. $n_t = 0$), $D\tilde{T}_a(\mathbf{a}', n)$ reduces to $[\tilde{\mathbf{A}}]$. Given that one eigenvalue of $\tilde{\mathbf{A}}$ is bigger than one, not all eigenvalues of $[\tilde{\mathbf{A}} - \mathbf{I}]$ will be negative. The magnitude of eigenvalues of $D\tilde{T}_b(\mathbf{b}', n)$ on the other hand depends on the correlation matrix $\mathbf{F}$ and might be bigger or smaller than one in absolute value. When all households are naive forecasters (i.e. $n_t = 1$), both $D\tilde{T}_a(\mathbf{a}', n)$ and $D\tilde{T}_b(\mathbf{b}', n)$ reduce to zero so that the associated eigenvalues are of course -1. Since the relevant eigenvalues switch signs I need to consider the steady state case $n = 0.5$ explicitly. In that case, the mapping becomes:

$$D\tilde{T}_a(\mathbf{a}', 0.5) = \left[(\mathbf{I} - {}^1\!/\!_2\tilde{\mathbf{A}})^{-1}\,{}^1\!/\!_2\tilde{\mathbf{A}}\right] = (2\mathbf{A}^{-1} - \mathbf{I})^{-1} \tag{78}$$

E-stability requires that

$$eig((2\tilde{\mathbf{A}}^{-1} - \mathbf{I})^{-1} - \mathbf{I}) < 0 \tag{79}$$

Where $eig(\mathbf{A})$ denotes the vector of eigenvalues λ associated with any matrix $\mathbf{A}$. The above is equivalent to

$$\Leftrightarrow eig((2\tilde{\mathbf{A}}^{-1} - \mathbf{I})^{-1}) - 1 < 0 \tag{80}$$

$$\Leftrightarrow 1 \oslash eig((2\tilde{\mathbf{A}}^{-1} - \mathbf{I})) - 1 < 0 \tag{81}$$

where $\oslash$ denotes the Hadamard division, i.e. elementwise division. Note further

$$\Leftrightarrow 1 \oslash (eig(2\tilde{\mathbf{A}}^{-1}) - 1) - 1 < 0 \tag{82}$$

$$\Leftrightarrow 1 \oslash (2 \oslash eig(\tilde{\mathbf{A}}) - 1) - 1 < 0 \tag{83}$$

$$\Leftrightarrow eig(\tilde{\mathbf{A}}) \oslash (2 - eig(\tilde{\mathbf{A}})) - 1 < 0 \tag{84}$$

however, since one of the eigenvalues of $\tilde{\mathbf{A}}$ is larger than one, this condition cannot be satisfied and steady state is not E-stable. □

A.7. Proof of Proposition VII

Proof. I guess that $\mathbf{z}^* = \bar{\mathbf{a}} = (\frac{1-\beta}{\kappa}\bar{\pi}, \bar{\pi}, 0)'$ (i.e. the rational expectations steady state) is a steady state. In the absence of structural shocks and given that $\mathbf{z}_t = \mathbf{z}_{t-1}$, neither forecast heuristic makes prediction errors so that $n_t = n = 0.5$ and the system becomes

$$y^* = (1-n)y^* + (1-n)c_{y,p}\hat{p}^* + ny^* - \frac{1}{\sigma}\left(\bar{\pi} + \phi_\pi(\pi_t - \bar{\pi} + \hat{p}^*) - (1-n)\pi^* - (1-n)c_{\pi,p}\hat{p}^* - n\pi^*\right) \tag{85}$$

$$\pi^* = \beta[(1-n)\pi^* + (1-n)c_{\pi,p}\hat{p}^* + n\pi^*] + \kappa y^* \tag{86}$$

$$\hat{p}^* = \pi^* - \bar{\pi} + \hat{p}^* \tag{87}$$

It is clear that $\pi^* = \bar{\pi}$, $y^* = \frac{1-\beta}{\kappa}\bar{\pi}$, and $\hat{p}^* = 0$ is a solution, thus confirming the initial guess. This is the case even if the weight of the random walk forecast is fixed at any value $n_t = n < \bar{n}$. However, due to the decay of past prediction errors and the forecast selection mechanism, only $n_t = n^* = 0.5$ can be a steady state. However, if $\bar{n} \leq 0.5$, there exists no fixed point $\bar{\tilde{\mathbf{a}}}$ so that the steady state does not exist. $\square$

A.8. Proof of Proposition VII

Proof. Suppose there exists a steady state at the zero lower bound where $\hat{p}_t = \hat{p}_{t-1} = \hat{p}^*$. In that case by the law of motion for $\hat{p}_t$ it must be that $\pi_t = \bar{\pi}$. However, a steady state with positive inflation and a constant price level gap is only possible if $\hat{p}^* = 0$, i.e. the steady state away from the ZLB. Therefore, there is no zero lower bound steady state under price level targeting.

$\square$

Appendix B. Estimation Results

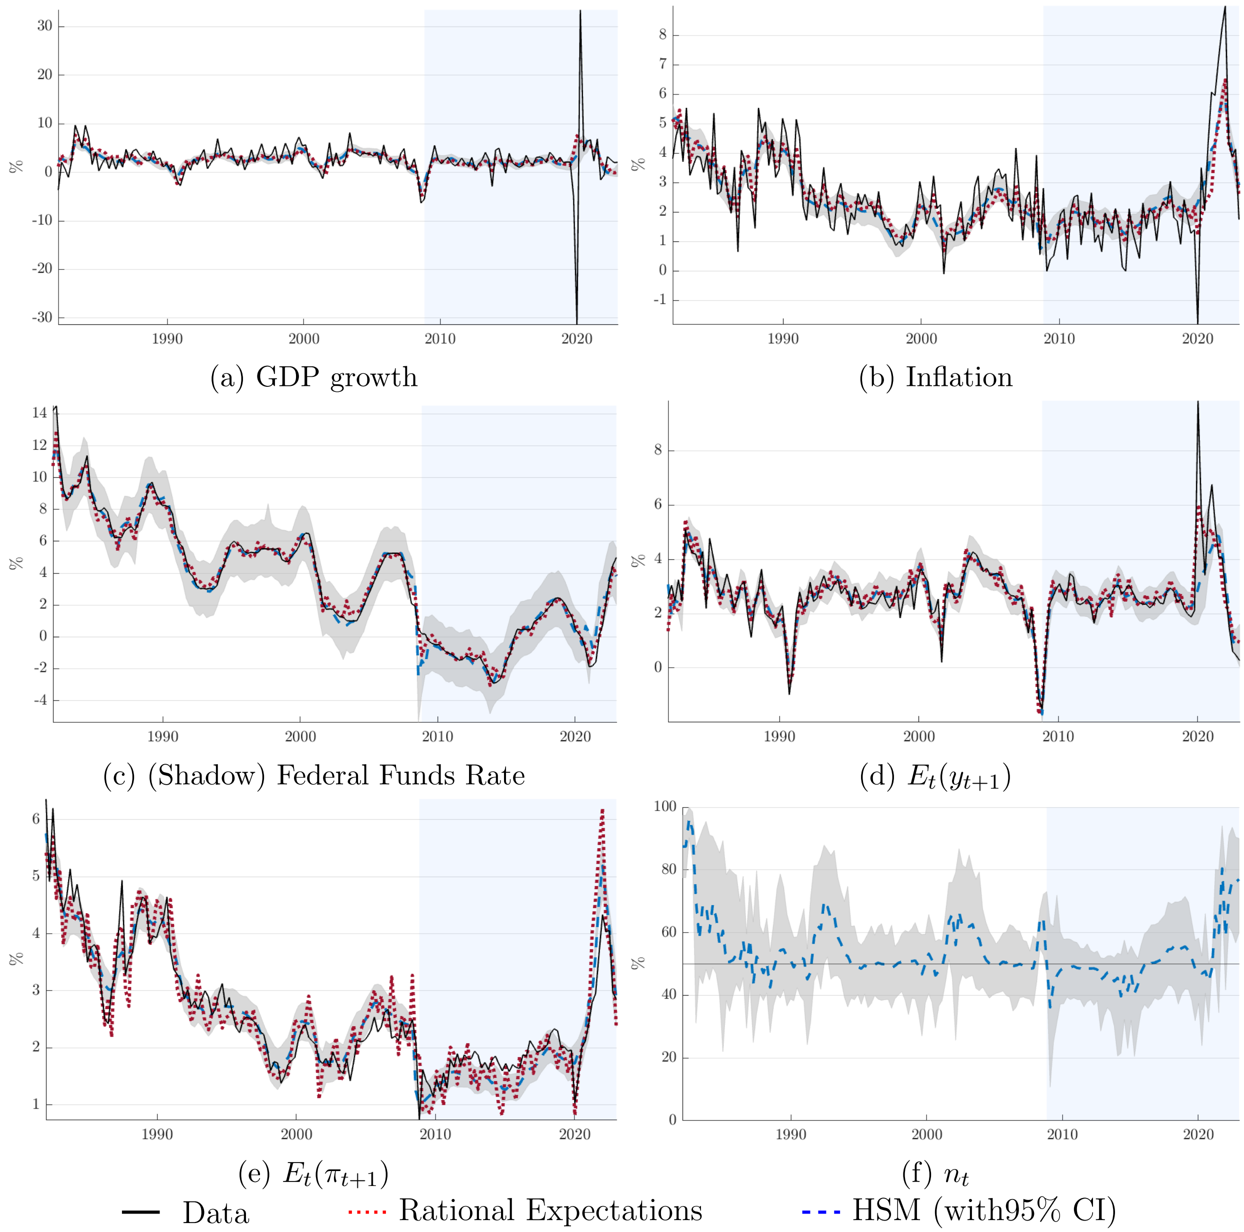

(a) GDP growth

(b) Inflation

(c) (Shadow) Federal Funds Rate

(d) $E_t(y_{t+1})$

(e) $E_t(\pi_{t+1})$

(f) n_t

Note: The panels displays the smoothed model states. The solid black line denotes the data; the blue dashed line shows the median model prediction of the heuristic switching model along with 95% confidence intervals in grey; the red dotted line depicts the median model prediction under rational expectations (RE). The blue-shaded area denotes the out-of-sample forecast period.

Figure B.1: Model Predictions

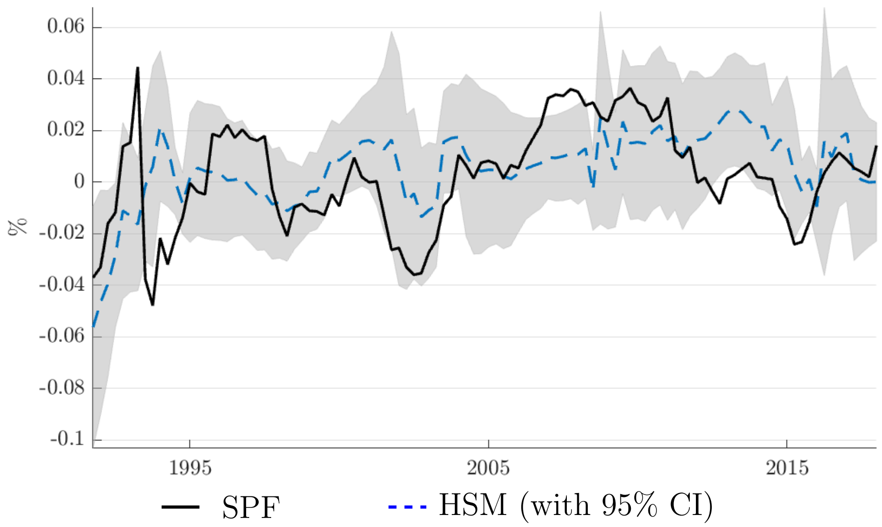

Note: This figure show the estimated time-varying sensitivity of smoothed long-run inflation expectations to inflation surprises. The blue dashed line depict the time-varying sensitivity of model-implied inflation expectations with 95% credible intervals shaded in grey. The black solid line depicts the time-varying sensitivity of median observed inflation expectations from the SPF. The blue-shaded area denotes the out-of-sample forecast period.

Figure B.2: Sensitivity of Implied Expectations to Inflation Surprises

Appendix C. Model Dynamics

C.1. Long-run Inflation Expectations & Monetary Policy

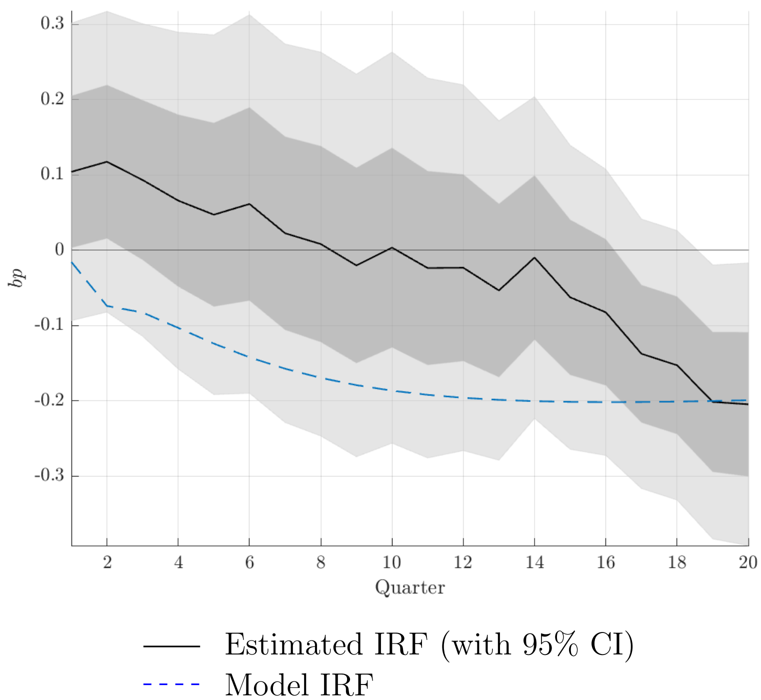

Note: This figure plots the estimated IRF of 10y inflation expectations to a 1pp monetary policy shock along with the equivalent model IRF. The data IRF is estimated using a local projection on a sample from 1990Q1:2007Q4 and uses Romer and Romer (2004) shocks to identify exogenous movements in monetary policy, which are used to instrument the 1y Treasury Rate. The model further includes lagged values of real GDP, CPI, the 1y Treasury Rate, and the dependent variable. Standard errors are heteroskedasticity and autocorrelation consistent with 12 lags. IRF is robust to 1) including the monetary policy shock directly instead of using it as an instrument; 2) using high frequency identified shocks (Gürkaynak et al., 2005).

Figure C.1: IRF of 10y Inflation Expectations to a 1pp Monetary Policy Shock

C.2. Forecast Errors

Angeletos et al. (2020) document two stylised facts of aggregate expectations: They initially under-react but later over-shoot the actual outcomes. I test whether my model of expectation formation fits these facts. I estimate the impulse response functions of average 1q ahead inflation forecast errors (i.e. $\pi_{t+1} - \hat{\mathbb{E}}_t \pi_{t+1}$) to a monetary policy shock. The left (right) panel of Figure C.2 shows the IRF of the forecast error after a positive (negative) monetary policy shock that decreases (increases) inflation. In both cases, we see an initial under-reaction of expectations. That is, after a positive (negative) monetary policy shock,the household initially expects higher (lower) inflation than eventually realises. However, after 5 periods, the forecast errors flip signs, i.e. they over-react. Thus, my model with heuristic switching forecasts fits the stylised facts of Angeletos et al. (2020). As a comparison, I also plot the forecast errors of the model under the assumption that the representative household puts no weight on the random walk forecasting function at all (i.e. $n_t = 0 \forall t$). In this case forecast errors are negative upon the impact but quickly return to zero, thus not corresponding to the the findings of Angeletos et al. (2020).

Figure C.2: IRF of Forecast Errors

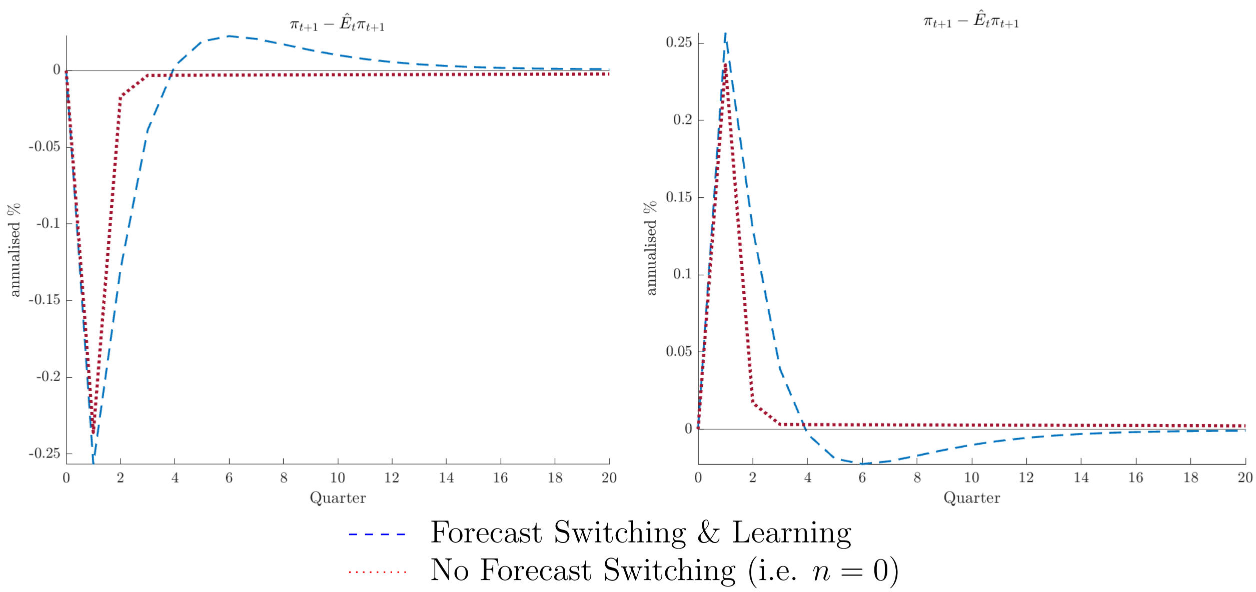

Note: This figure plots the IRF of the average inflation forecast errors to positive (right) and negative (left) 1pp monetary policy shock which occurs in period 1. The solid blue line depicts the forecast errors of the heuristic switching model, whereas the dotted blue line depicts the counterfactual forecast errors if the model contained adaptive learners only.

C.3. Model Moments

Figure C.3: Impulse Response Functions for various Degrees of De-anchoring

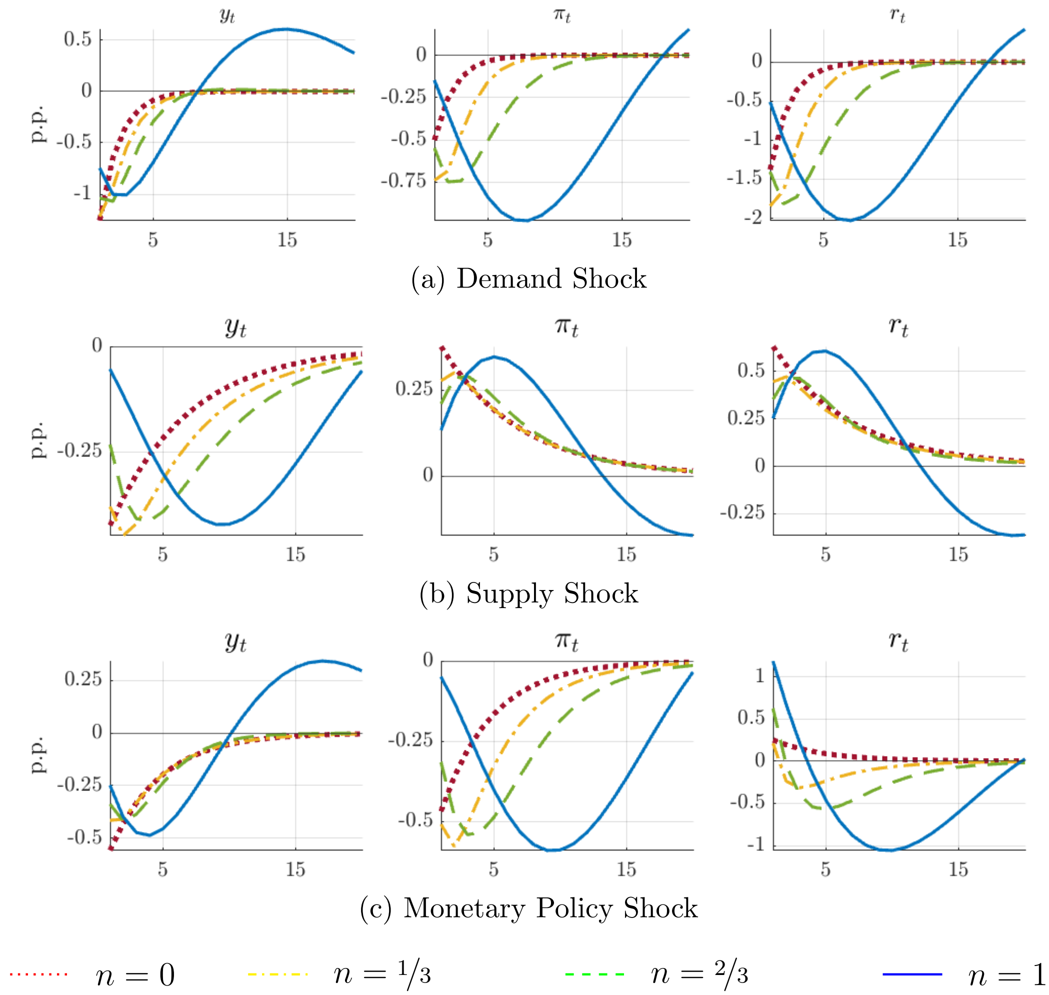

Note: This figure plots the IRF of output (left), inflation (middle), and interest rates (right) to a demand (upper panel), supply (middle panel), and monetary policy shock (lower panel) for different weights put on the random walk forecasting function.

Figure C.4: Model Variances & Covariance

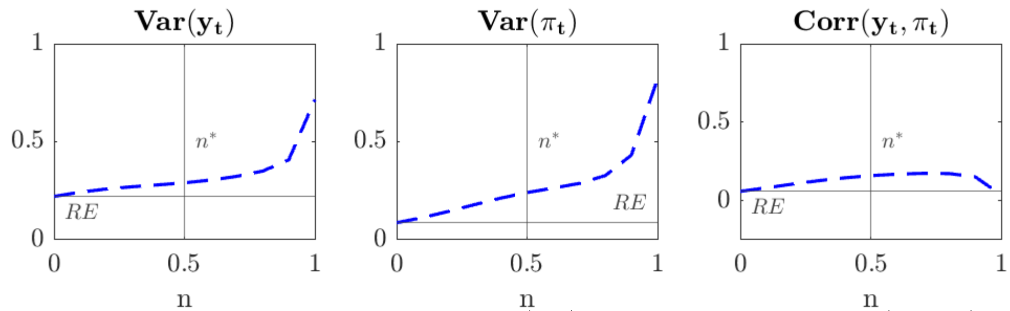

Note: This figure plots the variance of output (left), the variance of inflation (middle), and the covariance between the two (right) for different weights of the naive forecasting function.

C.4. Model Moments under Price Level Targeting

Figure C.5: Model Variances & Covariance under Price Level Targeting

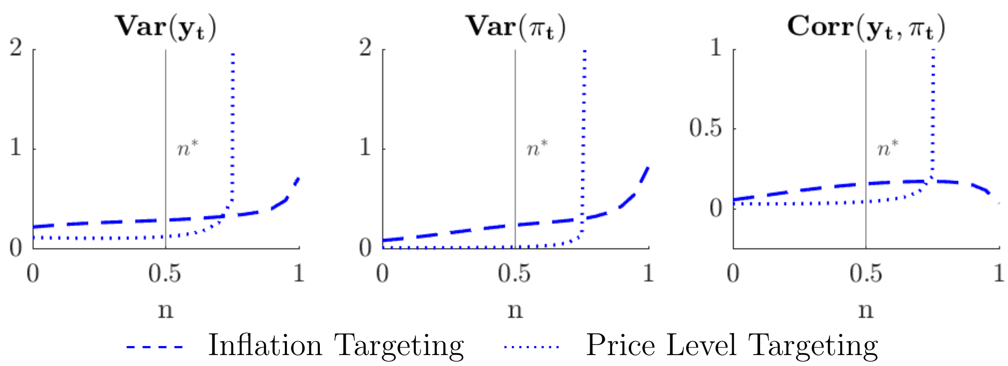

Note: This figure plots the variance of output (left), the variance of inflation (middle), and the covariance between the two (right) for different weights of the naive forecasting function und Inflation Targeting (dashed) and Price Level Targeting (dotted).

C.5. Counterfactual IRFs

Figure C.6: Counterfactual IRF Shutting Down one Updating Mechanism at a Time

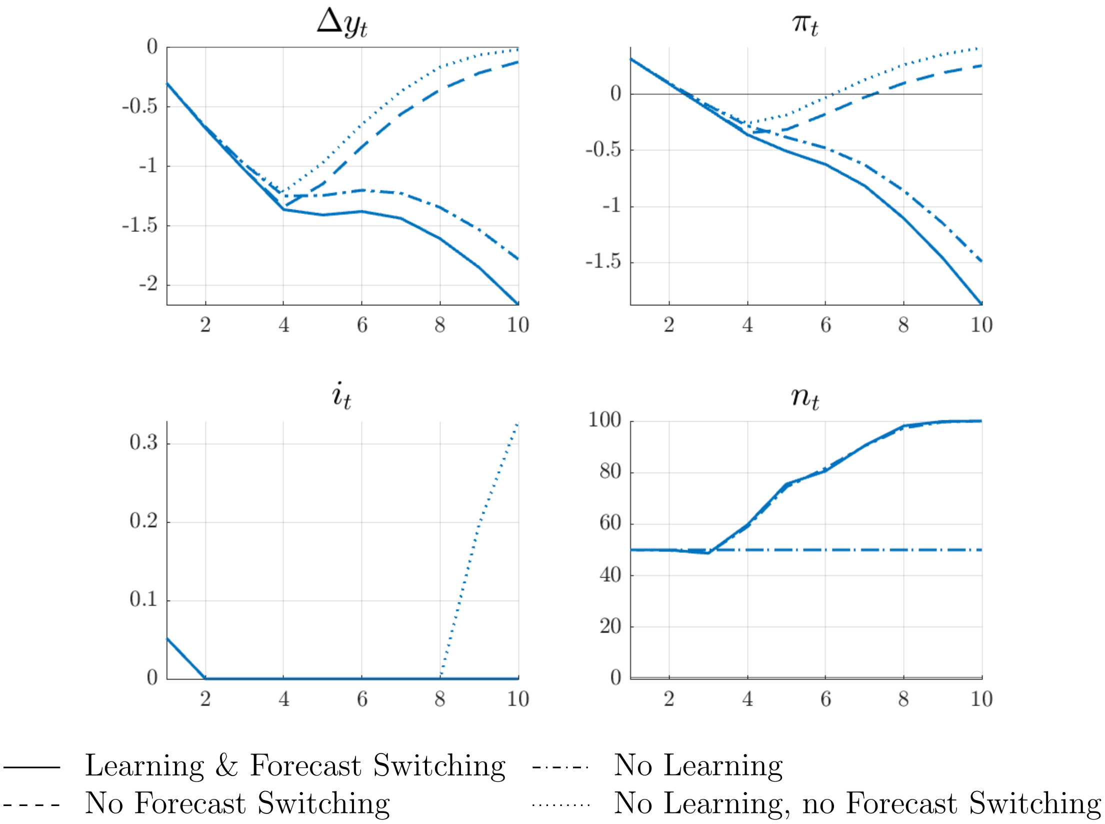

Note: This figure plots the IRF of output (left) and inflation (right) to four consecutive σ_ζ shocks to the discount factor under different assumptions about the expectation formation. The solid blue line shows the response of the full model (i.e. the same response as in Figure 6); the dashed lined shows the IRF of the model with a constant weight put on the random walk forecast (i.e. $n = 0$); the dotted lined shows the IRF of the model with a constant weight put on the random walk forecast (i.e. $n = 0.5$) and no updating of the adaptive forecasting function (i.e. $\Psi_t = \Psi$). The dash-dotted lined shows the IRF of the model with a time varying weight put on the random walk forecast but no updating of the adaptive forecasting function (i.e. $\Psi_t = \Psi$).

Appendix D. Extension: Responding to Long-term Inflation Expectations

In this section I investigate the performance of a Taylor Rule that also reacts to long-term inflation expectations, i.e. an implementation of the optimal policy rule by Gáti (2023).

$$r_t = \max\left[0, \bar{\pi} + \phi_\pi(\mathbb{E}_t\pi_{t+40} - \bar{\pi}) + \phi_\pi(\pi_t - \bar{\pi}) + \phi_y\hat{y}_t\right]$$

That is, the central bank reacts to deviations of long-term inflation expectations from its inflation target on top of its systematic response to output and inflation.

Table D.1: Policy Evaluation - Responding to Long-term Inflation Expectations

	$R^* = 6\%$		$R^* = 3\%$	
	IT	IT + long-run Expectations	IT	IT + long-run Expectations
ϕ_π^*	17.53	16.53	27.34	14.62
Welfare $\mathcal{W}$	-62.77	-62.77	-261.1	-259.41

Note: This table shows the optimal inflation response in a standard Taylor rule as well as a Taylor rule augmted with long-run inflation expectations based on a simulation of the non-linear model in high nominal interest rate environment (i.e. $R^* \sim 6\%$) and a low nominal interest rate environment (i.e. $R^* \sim 3\%$).

The simulation results in Table D.1 show that a Taylor rule that also responds to long-run inflation expectations can improve in welfare terms over Inflation Targeting. This difference is particularly pronounced in a low interest rate environment.

Appendix E. Extension: An Asymmetric Taylor Rule

In this section I investigate the performance of an asymmetric Taylor rule as in Bianchi et al. (2021) as an alternative policy rule under endogenously anchored expectations.

$$r_t = \max\left[0, \bar{\pi} + \mathbf{1}_{\pi_t<\bar{\pi}}\underline{\phi}_{\pi}(\pi_t - \bar{\pi}) + (1 - \mathbf{1}_{\pi_t<\bar{\pi}})\bar{\phi_{\pi}}(\pi_t - \bar{\pi}) + \phi_y \hat{y}_t\right]$$

That is, the central bank reacts more forcefully to inflation below target than inflation above target to address the asymmetric welfare loss implied expectations de-anchoring at the ZLB. I evaluate the performance of this rule in the same way as in Section 5.

Table E.1: Policy Framework Evaluation - Simulation Evidence

	$R^* = 6\%$			$R^* = 3\%$		
	IT	asym. IT	PLT	IT	asym. IT	PLT
Welfare $\mathcal{W}$	-126.48	-125.97	-126.59	-385.77	-416.88	-387.78
	[-192.42 -106.73]	[-197.41 -105.46]	[-213.51 -95.54]	[-566.56 -326.38]	[-113,988.11 -335.19]	[-566.46 -328.87]
ZLB Frequency (%)	4.69	3.25	5.75	30.75	25.38	24.38
	[0.88 12.63]	[1.97 12.03]	[9.50 20.50]	[14.50 37.31]	[12.56 32.88]	[13.63 37.56]
Risk of Defl. Spiral (%)	0.13	0.25	0.13	0.50	0.75	0.63
	[0.00 0.47]	[0.13 0.75]	[7.06 13.63]	[0.13 1.38]	[0.13 1.50]	[0.13 1.38]

Note: This table shows the simulation results under inflation targeting (IT) and price level targeting (PLT) in a high nominal interest rate environment (i.e. $R^* \sim 6\%$) and a low nominal interest rate environment (i.e. $R^* \sim 3\%$). Values in square brackets indicate 95% confidence bands, which are computed by taking 500 parameter draws from the draws of the MCMC sampler.

The simulation results in Table E.1 show that an asymmetric Taylor rule leads to very similar but slightly slower Welfare as symmetric Inflation Targeting in both high and low nominal interest rate environments. The reason is that, in the presence of supply shocks, the asymmetric response to inflation actually worsens the trade-off between output and inflation: in response to a positive cost-push shock that increases inflation and reduces output, the asymmetric policy rule becomes even more accommodative. However, in response to a positive cost push shock, less accommadative policy reduces the benefits in terms of higher growth