Forward Guidance with Unanchored Expectations*

Stefano Eusepi† Chris Gibbs‡ Bruce Preston§

Abstract

We study zero interest-rate policy in response to a large negative demand shock when long-run expectations can fall over time. Because falling expectations make monetary policy less effective by raising real interest rates, the optimal forward guidance policy makes large front-loaded promises to stabilize expectations. Policy is too stimulatory in the event of transitory shocks, but provides insurance against persistent shocks. The optimal policy is well-approximated by a constant calendar-based forward guidance, independent of the shock's realised persistence. The insurance property distinguishes our paper from other bounded rationality papers that solve the forward guidance puzzle and generates important quantitative differences.

Keywords: Optimal Monetary Policy, Learning Dynamics, Expectations Stabilization, Forward Guidance

JEL Codes: E32, D83, D84

*July 22, 2021. The authors thank seminar and conference participants at the Bank of Finland, SNDE Meetings 2019, Reserve Bank of Australia and the Virtual Australian Macroeconomics Seminar. The usual caveat applies. This research was supported by the Australian Research Council under the grant DP210101204.

†University of Texas, Austin. E-mail: stefano.eusepi@austin.utexas.edu.

‡University of Sydney. E-mail: christopher.gibbs@sydney.edu.au

§The University of Melbourne. E-mail: bruce.preston@unimelb.edu.au.

1 Introduction

This paper studies zero interest-rate policy in response to a large negative demand shock when long-run expectations can become unanchored and fall over time. Because falling inflation expectations make monetary policy less effective by raising real interest rates, the optimal forward guidance policy makes large front-loaded promises to stabilize expectations. The promises are strikingly different to a rational expectations analysis and far less effective at stabilizing the economy. We argue the predictions are more consistent with lived experience than those from rational expectations models, and that the results help organize thinking about practical policy concerns relating to unachored inflation expectations and the framing of policy announcements. The results also speak to contemporary debate on the merits of aggressive stimulus—front-loading forward guidance provides insurance against persistent shocks by limiting downward drift in expectations; but entails the cost of overheating the economy in the case of transitory shocks.

Figure 1 summarizes our central result. We determine the optimal forward guidance policy in Eggertsson and Woodford's (2003) classic thought experiment: a large negative demand shock, modeled as an unanticipated fall in the natural real interest rate, forces the central bank to choose a zero interest rate. With constant probability each period, the natural rate permanently reverts to steady state.[1] The first panel reports the optimal state-contingent promises attached to each realization of uncertainty. Unanchored inflation expectations require large front-loaded promises to stabilize the economy. In contrast, rational expectations require a less aggressive profile of promises that have a modest hump-shape. The second panel shows the total time at the zero lower bound for each belief assumption. Time at the zero lower bound rises more sharply with longer duration shocks under rational expectations. With unanchored expectations, the optimal state-contingent forward guidance policy is well-approximated by constant calendar-based forward guidance of four years (shown by the dashed line).[2] This is consistent with how many central banks implemented policy during the global pandemic.[3]

To derive these predictions we use the canonical New Keynesian model. As in a rational expectations analysis, we assume that agents understand the two-state-Markov process that describes the natural real rate. They observe the real rate; know the state of the economy at any time; and know the transition probabilities between states. When in the low state they also understand the central bank will always set interest rates to zero.

[1] The shock size and duration are identical to that studied in Eggertsson and Woodford (2003). As in their study, the shock is known to last at most 60 quarters. Other parametric assumptions are in section 5.

[2] The policy achieves a welfare outcome 95% of that delivered by optimal policy.

[3] For example, the Reserve Bank of Australia announced three years of forward guidance in March 2020.

Figure 1: Optimal Policy RE Commitment vs Unanchored Expectations

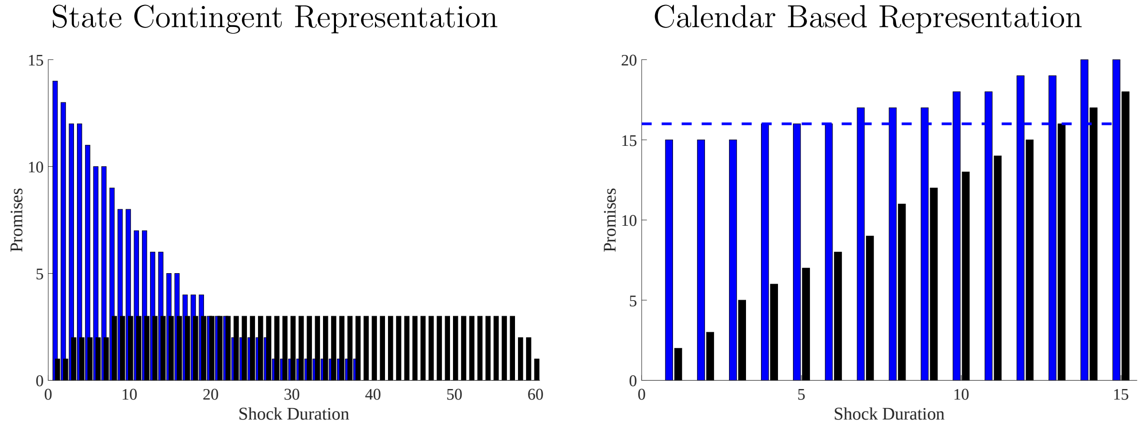

Notes: Rational expectations (black); unanchored expectations (blue). The dashed line in the right panel corresponds to a calendar based promise for holding the policy rate at zero for four years regardless of the duration of the shock during that period. .

We depart from rational expectations by allowing agents to be uncertain about medium and long-run developments in the economy. We assume that the forecast error generated by the onset of the negative natural rate shock causes agents to mark down their inflation and output gap beliefs. Because these mark downs affect subsequent dynamics and beliefs, the initial forecast error gets propagated over time. The basic idea is that agents have some understanding of the consequences of a large negative demand shock but are concerned the economic environment might have changed in ways they don't understand. They extrapolate from recent patterns in data to form expectations about longer term outcomes. How sensitive beliefs are to forecast errors provides a measure of how well anchored are expectations. The more sensitive to recent patterns in observed data, the more poorly anchored they are. How monetary policy manages the evolution of beliefs in the central question of the paper.

The analysis starts with a central bank that cannot credibly commit to forward guidance and must use short-term interest rate policy to manage aggregate demand. Unanchored expectations make the zero lower bound a more significant constraint on monetary policy: the economy spends substantially longer time at the zero lower bound when compared to a rational expectations economy. The difficulty is that inflation and output beliefs that are sufficiently pessimistic can cause the zero lower bound to be a constraint — even if the natural rate has reverted to steady state. Falling inflation expectations raise real interest rates and limits the ability of the central bank to normalize policy once the negative demand shock has passed. The optimal policy results in a period of below target inflation and slightly above target output, qualitatively similar to the experience of the United States in the aftermath of the Great Recession.[4]

Because unanchored expectations represent a constraint on conventional interest rate policy, we explore whether better equilibrium outcomes can be achieved by a credible commitment to future zero interest rate policy. To analyze forward guidance announcements with non-rational expectations, we make the following assumptions. The central bank is perfectly credible and agents perfectly understand the implications of forward guidance for the path of future interest rates. We assume that given the forward guidance announcement agents revise their beliefs about future interest rates, but leave unaltered their beliefs about future output and inflation. Agents do not evaluate the general equilibrium effects of the policy. This approach to implementing central bank communication follows Preston (2006) and Eusepi and Preston (2010), and is consistent with empirical evidence on the effects of forward guidance on survey data on professional forecaster expectations (see Crump, Eusepi, and Moench 2011, Del Negro, Giannoni, and Patterson 2012, Campbell, Evans, Fisher, and Justiniano 2012 and Andrade and Ferroni 2021). The approach is also an example of level k reasoning in which agents perform a single round of deductive reasoning in response to the policy announcement — see Farhi and Werning (2019).

We start with some simple analytics. We show that our model does not display a forward guidance puzzle: holding the natural rate at steady state, consider a central bank that lowers the short-term nominal policy rate to zero. As we increase the expected duration of zero interest rate policy the increase in output is bounded. In contrast, the output effects in the rational expectations model are unbounded. Of course, this paper is not the first to provide a resolution to the forward guidance puzzle. However, the analysis differs to many of these other papers because the effect of zero interest rate policy depends both on the steady state natural rate plus long-term inflation expectations. The endogeneity of long-term expectations to shocks and policy is a novel property of our model. As inflation expectations fall, zero interest rate policy becomes less effective.

Informed by these analytical results, we compute the optimal forward guidance policy in response to a large negative demand shock. As discussed, the promises are large and front-loaded. Optimal policy displays an insurance principle. The central bank makes large state-contingent promises for short-duration shocks to ensure inflation expectations don't fall in the event of a long duration shock. These aggressive announcements raise inflation and inflation expectations, which affords precious nominal space for interest rate policy, even after the period of forward guidance policy. The price of this insurance is a substantial rise in long-term inflation expectations that must be restrained by a contraction in real activity

[4] Whether the model can account for this period of economic history is an empirical matter. But the point is that a model of this kind can generate such patterns.

in the event the shock is short lived. And despite the aggressive nature of policy promises, optimal policy is associated with much larger fluctuations in inflation and output than under rational expectations analysis. Stabilization policy is more difficult.

The paper offers three applications of this result. The first compares our findings to other models which resolve the forward guidance puzzle. We show that many of these papers have dynamics that are nested by a modified representation of the canonical New Keynesian model. While our discussion gives emphasis to models proposed by Gabaix (2020) and Bilbiie (2018), the predictions of this general class of model are nearly indistinguishable from the optimal commitment rational expectations model. These models continue to have strong general equilibrium effects that are sufficient to generate substantial stimulus from forward guidance policy. That is, forward guidance about future interest rate policy has large effects on anticipated future outcomes for inflation and output that shift current equilibrium outcomes through income and substitution effects. In contrast, our model displays substantially weaker general equilibrium effects because beliefs about inflation and output respond only indirectly to forward guidance through its effect on recent outcomes for inflation and output and therefore beliefs.

The second application measures the costs of insurance. We consider a central bank that announces the optimal forward guidance policy at the time of the shock, but can renege on these promises if desirable. Because of the insurance principle which leads to substantial immediate stimulus, the central bank will in general renege on promised zero interest rate policy. Indeed, in the case of favorable short-duration shocks, raising interest rates early limits the extent to which higher inflation becomes entrenched in inflation expectations, requiring much less restraint in real activity. For this reason welfare is increased by about 20 per cent.[5]

The third application gauges the consequences of delay in the implementation of forward guidance policy. This is relevant to understanding different country experiences. For example, in the early 1990s Japan experienced a substantial downturn before much of the advance world adopted inflation targeting with the associated benefit of anchored long-term inflation expectations. Importantly, Japan did not employ unconventional policy until the early 2000s in form of quantitative easing, and forward guidance later still. In contrast, the United States confronted the challenges of zero interest policy during the global financial crisis and global pandemic having established a highly credible monetary policy with well anchored inflation expectations. The Federal Reserve was quick to implement forward guid-

[5] This thought experiment should not be interpreted as a policy proposal. Indeed, the reputational costs are likely to be substantial in low-interest rate environments in which forward guidance policies are likely to be required again. Rather, the experiment provides a measure of the costs of insurance.

ance and other unconventional monetary policies. We show that delay is costly because of the fall in inflation expectations. And the costs rise for more poorly anchored long-term expectations. As the implementation delay increases, the optimal promises are progressively more aggressive and flatter over the duration of the negative demand shock.

Related literature. Central to our result is the assumption that long-term expectations can become unanchored. The degree of anchoring is measured by the sensitivity of long-term expectations to short-term forecast errors or surprises. Using a structural model, Carvalho, Eusepi, Moench, and Preston (2019) provides evidence, for the United States and other countries including Japan, of a time-varying link between short-term inflation forecast errors and long-term forecast revisions using survey-based measures of expectations from professional forecasters. Using pass-through regressions of either macroeconomic news or movements in short-term expectations to long-term inflation expectations Gurkaynak, Levin, and Swanson (2010) and Beechey, Johannsen, and Levin (2011) show a similar link. In addition, Bern, Caselli, Giglioli, Gruss, and Lian (2018) provides a comprehensive measure of the degree of anchoring for a large set of countries and its impact on inflation dynamics.

Bounded rationality and imperfect information have been shown to have strong implications for monetary policy design. In Orphanides and Williams (2008), and other research in the learning literature, optimal monetary policy prescribes a more aggressive response to inflation and muted activism toward output gap stabilization compared to rational expectations. Closer to the current paper, Eusepi, Giannoni, and Preston (2019) show that in a model with drifting long-term beliefs monetary policy effectiveness is significantly hampered as the central bank has looser control of the term structure of interest rate expectations. Similarly, Gibbs and Kulish (2017) show that unanchored expectations complicates the design of disinflation policy and raises costs. These papers, however, do not study the zero lower bound and, more broadly, do not address forward guidance as an additional policy instrument. Preston (2006) and Eusepi and Preston (2010) discuss the importance of central bank communication in improving the monetary authority's stabilization trade-off. Similar to this paper, agents incorporate announced information about monetary policy rules into their interest-rate forecast, but they are unable to compute the general equilibrium effects on the economy. However, these papers do not address state contingent forward-guidance announcements at the zero lower bound.

Several papers address the monetary policy implications of the zero lower bound in new Keynesian models with bounded rationality and learning. Williams (2010) shows Taylor-type rules responding to the price level that are desirable under rational expectations do not perform as well when agents are boundedly rational and form expectations using constant gain estimators similar to what used here. In a similar model environment, various contributions study the response of economy to large shocks using global methods. Eusepi (2010), Evans, Guse, and Honkapohja (2008) and Evans and Honkapohja (2010) find that expectations drift pushes the economy into a prolonged period of low or negative inflation and low levels of output, very close to a liquidity trap equilibrium. While in these models central bank transparency about the policy rules improves outcomes, the central bank lacks forward guidance as a policy instrument. In addition, optimal policy is not analyzed.

Recent literature has analyzed the effects of forward guidance in models with bounded rationality and imperfect information. Farhi and Werning (2019), García-Schmidt and Woodford (2019), Woodford and Xie (2019) and Gabaix (2020) show that bounded rationality can eliminate the forward-guidance puzzle, making this policy instrument less powerful than under rational expectations. Angeletos and Lian (2018) and Wiederholt (2015) reach similar conclusions in an environment with information frictions. Imperfect common knowledge of forward guidance announcements can mute general equilibrium effects and therefore limit its impact on the economy. In fact, Angeletos and Lian (2018) show that such frictions lead to the same model representation as what is obtained under bounded rationality. Similar to our paper, they show that front-loading of fiscal policy announcement is desirable, although the mechanism is different. As shown in section 6, general equilibrium effects can still be quite powerful in these model environments, making optimal policy quite close to the full information benchmark. In addition, these papers do not address the implications from dynamic adjustment of beliefs as a result of un-anchoring, which is central to our analysis.

Lastly, in Andrade, Gaballo, Mengus, and Mojon (2019) and Bodenstein, Hebden, and Winkler (2019) forward guidance is shown to have limited power because of either limited credibility of the central bank, or imperfect information about its preferences. In this paper we make the stark assumption that the central bank is fully credible, so the lower effectiveness is sourced entirely to weaker general equilibrium effects induced by bounded rationality and learning.

2 The Model

We use a simple New Keynesian model to study zero interest rate policy. This facilitates analytical results and comparison to other recent papers on this topic. Some assumptions are made for expositional simplicity — for example linear disutility of labor supply. Further details on the microfoundations can be found in Woodford (2003) and Gali (2008).

2.1 Optimal Decisions

A continuum of households i on the unit interval maximize utility

$$\hat{E}_t^i \sum_{T=t}^{\infty} \bar{C}_T \beta^{T-t} \left[\frac{(c_T(i))^{1-1/\sigma}}{(1-1/\sigma)} - \chi n_T(i) \right],$$

where $0 < \beta < 1$, $\sigma > 0$ and $\chi > 0$, by choice of sequences for consumption, $c_t(i)$, and labor supply, $n_t(i)$, subject to the flow budget constraint

$$c_t(i) + b_t(i) \leq (1 + R_{t-1}) \pi_t^{-1} b_{t-1}(i) + W_t n_t(i)/P_t + \Gamma_t(i)/P_t$$

and the No-Ponzi condition

$$\lim_{T\to\infty} \hat{E}_t^i \left(\prod_{s=0}^{T-t} (1 + R_{t+s}) \pi_{t+s+1}^{-1} \right)^{-1} b_{T+1}(i) \geq 0.$$

The variable $b_t(i) \equiv B_t(i)/P_t$ denotes real bond holdings (which in equilibrium are in zero net supply), R_t the nominal interest rate, $\pi_t \equiv P_t/P_{t-1}$ the inflation rate, W_t is the hourly wage, $\Gamma_t(i)$ dividends from equity holdings of firms and $\bar{C}_T$ exogenous preference shifter. The operator $\hat{E}_t^i$ denotes subjective expectations, which might differ from rational expectations.

A continuum of monopolistically competitive firms maximize profits

$$\hat{E}_t^j \sum_{T=t}^{\infty} \xi^{T-t} Q_{t,T} \left[p_t(j) y_T(j) - W_T n_T(j) \right]$$

by choice of $p_t(j)$ subject to the production technology and demand function $y_T(j) = n_T(j) = (p_t(j)/P_T)^{-\theta} y_T$ for all $T \geq t$, with the elasticity of demand across differentiated goods an exogenous process satisfying $\theta > 1$; and exogenous probability $0 < \xi < 1$ of not being able to reset their price in any subsequent period. When setting prices in period t, firms are assumed to value future streams of income at the marginal value of aggregate income in terms of the marginal value of an additional unit of aggregate income today giving the stochastic discount factor $Q_{t,T} = \beta^{T-t} \left[(P_t y_t)/(P_T y_T) \right]^{(1/\sigma)}$.

For any beliefs satisfying standard probability laws, to a first-order log-linear approximation in the neighborhood of a zero-inflation steady state, optimal individual consumption and pricing decisions can be expressed as

$$\hat{c}_t(i) = \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta)\hat{w}_T - \beta\sigma \left(\hat{R}_T - \hat{\pi}_{T+1} - (\bar{c}_T - \bar{c}_{T+1}) \right) \right] \quad (1)$$

$$\hat{p}_t(j) = \hat{E}_t^j \sum_{T=t}^{\infty} (\xi\beta)^{T-t} \left[(1-\xi\beta)\hat{w}_T + \xi\beta\hat{\pi}_{T+1} \right] \quad (2)$$

where for any variable z_t, $\hat{z}_t = \ln(z_t/\bar{z})$ the log-deviation from steady state $\bar{z}$, with the exceptions $\hat{p}_t(j) = \ln(p_t(j)/P_t)$, $\hat{R}_t = \ln\left[(1+R_t)/\left(1+\bar{R}\right)\right]$, and $\bar{c}_t = \ln\left(\bar{C}_t/\bar{C}\right)$. With a slight abuse of notation, the caret denoting log deviation from steady state is dropped for the remainder, so long as no confusion results.

In a symmetric equilibrium $c_t(i) = c_t = w_t \equiv W_t/P_t = n_t = y_t$ for all i, $p_t(j) = p_t(j)$ and $b_t(i) = b_t(j) = 0$ for all i, j.[6] Aggregating across the continuum of households and firms, and imposing market-clearing conditions, the economy is described by the aggregate demand and supply equations

$$x_t = \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \sigma \left(R_T - \pi_{T+1} - r_T^n \right) \right] \quad (3)$$

$$\pi_t = \hat{E}_t \sum_{T=t}^{\infty} (\xi\beta)^{T-t} \left[\kappa x_T + (1-\xi)\beta\pi_{T+1} \right] \quad (4)$$

where the output gap is defined as

$$x_t = y_t - y_t^n = w_t$$

the difference between output and the natural rate of output, the level of output determined by a flexible price economy: here $y_t^n = 0$. We assume agents understand this equilibrium relationship between wages and the output gap. This is without loss of generality. The associated natural rate of interest $r_t^n = (\bar{c}_t - \hat{E}_t\bar{c}_{t+1})$ is determined by fluctuations in the propensity to consume, an exogenous process to be discussed. Average beliefs are defined as

$$\int_0^1 \hat{E}_t^i di = \int_0^1 \hat{E}_t^j dj = \hat{E}_t.$$

The aggregate demand equation determines the output gap as the discounted expected value of future wages, with the second term capturing variations in the real interest rate, applied in future periods, due to changes in the nominal interest rate and goods price infla-

[6] The optimal consumption decision rule implictly made use of this equilibrium requirement.

tion. That expected future dividends are irrelevant to consumption plans, to the first-order, reflects the assumption of an infinite Frisch elasticity of labor supply. The aggregate supply curve determines inflation as the discounted future sequence of marginal costs and the inflation rate. The slope of the Phillips curve is measured by $\kappa = (1 - \xi\beta)(1 - \xi)/\xi$.

2.2 Uncertainty and Information

We use a standard thought experiment to characterize uncertainty. The natural rate of interest is unexpectedly negative in period 1 and reverts to its steady-state value with constant probability. When the natural rate is negative, monetary policy is constrained by the zero lower bound on nominal interest rates. Call this the low state, L. The nominal interest rate remains at zero while the economy is in the low state. When the natural rate is at steady state we call this the high state, H. The high state is absorbing.

Monetary policy. On return to the high state, nominal interest rates are determined by central bank policy. The central bank has two policy instruments at its disposal. The short-term nominal interest rate, which becomes available only after the shock reverts to the high state; and forward guidance, which takes the form of a policy announcement at the time the economy switches to the low state, detailing the state-contingent path for the nominal interest rate during and after the ZLB episode. Forward guidance announcements are *fully credible* and fully understood by private agents. We discuss this policy in section 5.

Information. Agents are boundedly rational and have imperfect information about the true data-generating processes of key equilibrium variables they need to forecast to make spending and pricing decisions. Before describing how their expectations are formed, let us clarify what private agents know and do not know in this economy. Agents perfectly observe the exogenous natural rate of interest, r_t^n, and understand its two-state process, including the constant probability of switching. At all times they are aware of the current state of the economy. In addition, agents correctly anticipate that the nominal rate will remain at the ZLB while the economy remains at the low state.

These assumptions are present in a full information rational expectations analysis. We depart from a rational expectations analysis by assuming that agents are uncertain about the equilibrium level of output and inflation in both the high and low state and revise their views about the economy in response to observed data. Agents are also uncertain about the evolution of the nominal interest rate in the high state. However, in forming their expectations they make use of two pieces of information. First, expectations are shaped by forward guidance announcements. Second, agents use the Fisher equation to incorporate their knowledge about the natural rate into their nominal interest rate forecast. As a result, inflation expectations are the key driver of interest rate expectations.

Taken together, our modeling assumptions suppose that agents have some understanding of the consequences of a large negative demand shock, but form expectations on the basis that the economic environment may have changed in ways they don't understand. Agents extrapolate from patterns observed in recent data and the extent to which expectations are sensitive to these patterns a measure of how well anchored expectations are.

2.3 The case of constrained policy

We first consider an economy where the central bank cannot use credible announcements to influence expectations. This is a valuable benchmark to introduce the model and discuss key properties relative to a full information rational expectations economy.

Subjective expectations. Given the information available to them, agents form estimates of equilibrium macroeconomic variables like econometricians. Expectations are based on the forecasting model

$$z_t = z_S + \bar{\omega}_t + e_t \tag{5}$$

$$\bar{\omega}_{t+1} = \rho\bar{\omega}_t + u_{t+1} \tag{6}$$

where $S \in [H, L]$ and

$$z_t = \begin{bmatrix} \pi_t \\ x_t \end{bmatrix}, \quad z_S = \begin{bmatrix} \pi_S \\ x_S \end{bmatrix}, \text{ and } \bar{\omega}_t = \begin{bmatrix} \bar{\omega}_t^{\pi} \\ \bar{\omega}_t^{x} \end{bmatrix}$$

and $0 \leq \rho \leq 1$ a parameter; e_t and u_t i.i.d. with $J = E\left[e_t e_t'\right]$ and $Q = E\left[u_t u_t'\right]$. These priors are invariant to the state.

The forecasting model has two components. The first is a state-dependent constant, reflecting agents' priors about the equilibrium outcomes in each state. This permits agents to have some understanding of the economic consequences of a large demand shock. This component of expectations is exogenous to the model and independent of monetary policy. The second component captures uncertainty about the level of output and inflation in the medium to long term. This term, which we call 'a drift', can capture different degrees of expectations anchoring in response to the economy entering the low state. This component of expectations is endogenous to the model and influenced by monetary policy. The larger is the parameter ρ, the more persistent are the effects of the shock on expectations. In the special case $\rho = 1$ we have an example of a shifting end-point model in the language of Kozicki and Tinsley (2001). For sufficiently large values of ρ, the negative state affects expectations beyond the horizon the shock, as market participants anticipate the recession will deliver a prolonged period of below-target inflation and negative output gaps.

Define $\omega_{t+1|t} \equiv \hat{E}_t\left(\bar{\omega}_{t+1}|z_t, z_{t-1}, \omega_t\right)$ to be the estimate of the unobserved drift. We describe how this estimate is formed below. Conditional on being in the low state, in any period t, expectations of inflation and output are given by

$$\begin{aligned}\hat{E}_t\left(x_{t+1}|S=L\right) &= (1-\delta)(\omega^x_{t|t-1}+x_L)+\delta(\omega^x_{t|t-1}+x_H)\\ \hat{E}_t\left(\pi_{t+1}|S=L\right) &= (1-\delta)(\omega^\pi_{t|t-1}+\pi_L)+\delta(\omega^\pi_{t|t-1}+\pi_H)\end{aligned}$$

where $\omega^x_{t|t-1}$ and $\omega^\pi_{t|t-1}$ are the beliefs about the current level of output and inflation. In absence of forward guidance, beliefs about interest rates satisfy

$$\begin{aligned}\hat{E}_t\left(R_{t+1}|S=L\right) &= (1-\delta)\times 0+\delta(\omega^\pi_{t|t-1}+r_H)\\ \hat{E}_t\left(r_{t+1}|S=L\right) &= (1-\delta)r_L+\delta r_H\end{aligned}$$

where $r_L < 0 < r_H$ are the values of the natural rate in the low and high states. The conditional expectation of next-period interest rates uses the fact that agents know interest rates will be zero in the low state; and that they form expectations of nominal interest rates in the high state using the Fisher equation to determine the nominal interest rate drift as

$$\omega^R_{t|t-1} = \omega^\pi_{t|t-1} + r_H.$$

Given this information structure the expectation of any variable at any future date, conditional on the low state in period t, can be computed as follows. Taking output as an example, we have

$$\begin{aligned}\hat{E}_t x_{T+1} &= (1-\delta)^{T+1-t}(x_L+\rho^{T-t}\omega^x_{t|t-1})+\delta\sum_{j=t}^{T}(1-\delta)^{j-t}(x_H+\rho^{T-t}\omega^x_{t|t-1}) &(7)\\ &= (1-\delta)^{T+1-t}\,x_L+\left(1-(1-\delta)^{T+1-t}\right)x_H+\rho^{T-t}\omega^x_{t|t-1} &(8)\end{aligned}$$

using the property

$$\hat{E}_t\omega^x_{T|t-1} = \rho^{T-t}\omega^x_{t|t-1}$$

for $T > t$. In the first line, the first term captures the expected value of output conditional on remaining in the low state for the next $T+1$ periods, which occurs with probability $(1-\delta)^{T+1}$. The second term represents the expectation of output conditional on returning to the high state, which occurs with probability

$$\delta \sum_{j=t}^{T} (1-\delta)^j = 1 - (1-\delta)^{T+1}.$$

In the second line, the final term captures impact of perceived drifts on forecasts, an effect independent of the state. For sufficiently high values of ρ, changes in the estimated drift have impact on forecast horizons beyond the low state. We construct expectations of the remaining variables in similar fashion.

A Special Case. When agents' beliefs satisfy

$$\omega^x_{t|t-1} = \omega^\pi_{t|t-1} = 0$$

they have rational expectations. Standard calculations demonstrate there is a rational expectations equilibrium with $x^H = \pi^H = 0$ and $R^H = r^H$ in the high state and

$$x_L = \frac{\sigma(\beta(\delta-1)+1)}{\beta(\delta-1)\delta + (\delta-1)\kappa\sigma + \delta} r_L \quad (9)$$

$$\pi_L = \frac{\kappa\sigma}{\beta(\delta-1)\delta + (\delta-1)\kappa\sigma + \delta} r_L \quad (10)$$

and $R_L = 0$ in the low state. A large decline in the natural rate of interest creates an aggregate demand shortfall and depresses inflation below the inflation target. Formally, this solution corresponds to the optimal policy under discretion.[7]

Throughout the paper we endow agents with some knowledge of the rational expectations equilibrium outcomes in response to a large negative demand shock. We assume agents' prior beliefs about the equilibrium outcomes in each state coincide with the rational expectations optimal discretion equilibrium. An alternative approach would be to model learning about the state-specific equilibrium values of inflation and output. However, the dynamics of the economy would then be sensitive to initial beliefs. For example, if at the time of the negative shock the initial belief was for a mild recession, the equilibrium drop in output gap and inflation would be higher than under full information. Conversely, if initial beliefs were pessimistic the equilibrium outcomes would be lower. We chose then to center beliefs at rational expectations. The impact response in the two economies will then be identical. The effects of learning dynamics are then relative to this well understood rational expectations benchmark, providing a clear context for our results.

[7] A unique bounded rational expectations solutions requires the denominator in the output and inflation expressions to be positive. The technical appendix provides derivations.

Objective Beliefs. Using the above assumptions to evaluate expectations in (3) and (4), provides the true data-generating process

$$x_t^H = -\sigma(R_t - r_H) + \frac{1-\beta}{1-\beta\rho}\omega_{t|t-1}^x + \frac{\sigma(1-\beta)}{1-\beta\rho}\omega_{t|t-1}^\pi \tag{11}$$

$$\pi_t^H - \kappa x_t^H = \frac{(1-\alpha)\beta}{1-\alpha\beta\rho}\omega_{t|t-1}^\pi + \frac{\kappa\alpha\beta}{1-\alpha\beta\rho}\omega_{t|t-1}^x \tag{12}$$

in the high state and

$$x_t^L = x_L + \frac{1-\beta}{1-\beta\rho}\omega_{t|t-1}^x + \sigma\left[\frac{1-\beta}{1-\beta\rho} + \frac{\beta(1-\delta)}{1-\beta\rho(1-\delta)}\right]\omega_{t-1}^\pi \tag{13}$$

$$\pi_t^L - \kappa x_t^L = \pi_L - \kappa x_L + \frac{(1-\alpha)\beta}{1-\alpha\beta\rho}\omega_{t|t-1}^\pi + \frac{\kappa\alpha\beta}{1-\alpha\beta\rho}\omega_{t|t-1}^x \tag{14}$$

in the low state. Besides the state-dependent constants, the structure of the aggregate supply curve is invariant across the low and high states. The perceived output gap and inflation drifts affect inflation in standard ways: higher inflation and output gap expectations raise equilibrium inflation through strategic complementarity in price setting and rising marginal costs. The size of these effects depends on the persistence of beliefs. The greater the persistence, the greater the effect of a shift in period t beliefs.

In contrast to aggregate supply, the aggregate demand equation differs across states. While the wealth effects from anticipated future income are identical (the terms in $\omega_{t|t-1}^x$), the substitution effects from interest rate policy differ (the terms in $\omega_{t|t-1}^\pi$). Consider output in the high state. Monetary policy affects demand through the contemporaneous interest rate relative to the natural rate and also through expectations of future real interest rates, given by the final term. These two effects are different in the low state. The contemporaneous effect of interest rates is given by $-\sigma(R_t - r_H)$ in the high state and by $-\sigma(0 - r_L)$ in the low state. Because of the zero lower bound, this generates a decline in demand (r_L being negative). This decline is captured by the constant term x_L in equation (9), representing the low state equilibrium output gap that would obtain under rational expectations. The elasticity with respect to expected future real rates is also substantially larger in the low state. Falling inflation expectations, which lead to rising real rate projections through nominal rigidities, have much bigger contractionary effect. In the low state, the elasticity is larger by the amount

$$\sigma\frac{\beta(1-\delta)}{1-\beta\rho(1-\delta)}.$$

The term multiplying σ is the expected discounted duration of the shock: the numerator being the discounted probability of being in the low state next period; and one over the denominator the average duration of the shock adjusting for the effects from discounting and the persistence of beliefs. The more persistent are beliefs the larger the elasticity — expectations that are poorly anchored and drift downward have a highly contractionary effect on aggregate demand.

Belief Updating. Estimates of the inflation and output gap drifts are obtained from the model described in equations (5) and (6). Following Sargent and Williams (2005) we study a steady state Kalman filter under the assumption that the variance-covariance matrix of the drifts' innovations is proportional to the variance-covariance matrix of short term innovations: that is $Q = \hat{c}J$ for a scalar $\hat{c}$. The appendix shows that updating of beliefs satisfies

$$\omega^x_{t+1|t} = \rho\omega^x_{t|t-1} + \rho\gamma(x_t - x_S - \omega^x_{t|t-1}) \quad (15)$$

$$\omega^\pi_{t+1|t} = \rho\omega^\pi_{t|t-1} + \rho\gamma(\pi_t - \pi_S - \omega^\pi_{t|t-1}) \quad (16)$$

where $0 < \gamma < 1$ is a function of the parameters ρ and $\hat{c}$. The learning gain is then $g \equiv \rho\gamma$. We use the notation $\omega_{t|t-1}$ to emphasize the beliefs about output and inflation period t are formed using data available to period $t-1$.[8]

Forecast Errors and Initializing Beliefs. Substituting the expressions for equilibrium output and inflation, (13) and (14), into the belief updating equations, (15) and (16), gives the law of motion for beliefs. Collectively these equations give the state-space representation of the model, the true data generating process. Forecast errors and, therefore, beliefs are in general independent of the normalizing constants $\{x_S, \pi_S\}$. For this reason the precise choice of normalizing constant does not matter for our results.[9] But this raises the question: what is the source of forecast error? Stated differently, what is the source of the initial shift in beliefs? Here the normalizing constants play an important role.

The natural rate shock is unexpected and occurs in period 1. Beliefs at the time of shock are inherited from period zero and assumed to be consistent with the steady state:

$$E_0 x_1 = x_H + \omega^x_{1|0} = x_H$$

$$E_0 \pi_1 = \pi_H + \omega^\pi_{1|0} = \pi_H$$

since $\omega^x_{1|0} = \omega^\pi_{1|0} = 0$. After the shock in period 1, we assume agents compute the forecast

[8] This avoids a complex simultaneity arising from having current beliefs depend on current macroeconomic outcomes.

[9] The normalizing constants clearly affect the trajectories of output and inflation through (13) and (14). But they do not affect the endogenous evolution of beliefs.

errors for period 1 as

$$x_1 - E_0 x_1 = x_L - x_H$$
$$\pi_1 - E_0 \pi_1 = \pi_L - \pi_H.$$

The period zero conditional expectations are ‘sticky’ in the sense that even though agents understand the natural rate shock gives rise to equilibrium outcomes $\{x_L, \pi_L\}$ they do not use this knowledge to infer their true level of surprise *given they observe the natural rate shock* r_L.[10] Beliefs in period 2 are determined as

$$\omega^x_{2|1} = g(x_L - x_H)$$
$$\omega^\pi_{2|1} = g(\pi_L - \pi_H).$$

The normalizing constants therefore determine the size of the initial forecast error. By choosing these constants to be given by the discretion equilibrium under rational expectations we make the impact effect of the shock identical under each belief assumption. Moreover, subsequent equilibrium outcomes differ only because of the endogenous propagation of this forecast error. Beliefs in periods $t > 2$ then satisfy (15) and (16).

Of course, this is not an optimal filtering problem. But all that matters for our analysis is that agents make some forecast error when the low state is realized. And given an initial forecast error, the interesting economic questions concern how different monetary policies propagate this forecast error over time. An alternative interpretation is this is a pessimism shock.

Extrapolation bias. The time-varying drift encodes uncertainty about medium-term developments in the economy. Extrapolation bias is a general equilibrium property of the model: subjective beliefs are more persistent than objective beliefs. Eusepi, Giannoni, and Preston (2019) show this information friction is consistent with the behavior of professional survey forecast data for the US over the period 1960-2007, when the economy was away from the zero lower bound. Furthermore, this mechanism is consistent with growing evidence that observed measures of expectations exhibit extrapolation bias (Fuster, Laibson, and Mendel 2010, Bordalo, Gennaioli, Ma, and Shleifer 2018 and Angeletos, Huo, and Sastry 2020). While it always amplifies and propagates shocks to the natural rate of interest relative to a rational expectations analysis, we show below that this mechanism produces stronger

[10] This is a standard assumption in the adaptive learning literature. While economic decisions assume agents observe and respond to shocks, their forecasting models are mis-specified because they are assumed not to make use of this information. Alternatively, we could assume the low state is unobserved by the agents. This would introduce a meaningful signal extraction problem but would add a layer of complexity to the analysis and cloud comparison with the rational expectation benchmark.

equilibrium effects when the economy reaches the zero lower bound. Uncertainty about the dynamics of output and inflation have general equilibrium effects that create a drag on economic recovery after a large negative shock. We show that when economic developments depress long-term expectations there are striking implications for monetary policy relative to both rational expectations and other models of bounded rationality.

3 Optimal Policy without Forward Guidance

This section provides benchmark results on optimal policy at the zero lower bound. We show that learning creates challenges for stabilization policy. Downward drift in beliefs can result in the zero lower bound being a constraint on policy even when the natural rate has reverted to the high state. In general the economy will experience longer durations of zero interest rate policy relative to rational expectations, with protracted periods of inflation and output below target.

3.1 The Policy Problem

Following Eggertsson and Woodford (2003) the central bank minimizes the loss function

$$E_t L_t = \sum_{T=t}^{\infty} \beta^{T-t} \left(\pi_T^2 + \lambda x_T^2 \right) \tag{17}$$

where $0 < \beta < 1$ and $\lambda > 0$ determines the relative weight placed on inflation stabilization versus output gap stabilization. This is the welfare-theoretic loss function implied by the microfoundations under both rational expectations and learning. The central bank has rational expectations and knows the true data-generating process.

The optimal policy problem minimizes this loss subject to the constraints implied by private behavior and the shock process for the natural rate. The appendix writes down the optimal policy problem, first-order conditions and describes the solution algorithm. The solution to the policy problem has the following characteristics. In response to a large negative shock to the natural rate, nominal interest rates become constrained by the zero lower bound. The optimal policy response has three regimes. While the shock persists it is optimal for the central bank to maintain a zero interest rate policy. During this time, the dynamics of the economy are given by (13) and (14). This is the first regime. When the shock reverts to its steady state value, interest rate policy must satisfy the constraints (11) and (12). Whether it is desirable to raise interest rates from zero depends on beliefs. If inflation and output expectations are sufficiently pessimistic and negative, then it will be optimal to maintain the zero interest rate policy. Output dynamics are then given by

$$x_t^H = \sigma r_H + \frac{1-\beta}{1-\beta\rho}\omega^x_{t|t-1} + \frac{\sigma(1-\beta)}{1-\beta\rho}\omega^\pi_{t|t-1}$$

where σr_H measures the effective stimulus from interest rate policy. This is the second regime. As beliefs recover, it eventually becomes desirable to raise interest rates with dynamics satisfying (11) and (12). This is the third regime. A fundamental difference from rational expectations analysis under discretion then is the presence of the second policy regime in the optimal policy problem. This regime is a consequence of extrapolation bias.

3.2 The Experiment

Because of the zero lower bound constraint, the optimal policy problem is non-linear. We therefore provide a numerical characterization of optimal policy. The thought experiment assumes the natural rate of interest is unexpectedly negative in period 1 taking a value of -2 percent per annum. The natural rate reverts back to the steady state value of $r_H > 0$ with probability 0.1 in each period. The steady state value of the natural real rate is assumed to be 4 percent per annum.[11]

The natural rate shock reverts to steady state with certainty after 60 periods. To construct a solution in forward-looking models requires working recursively backwards in time from a known terminal point. Eggertsson and Woodford (2003) find that this truncation point permits a good approximation of the two-state-Markov process for the natural rate in the rational expectations optimal commitment solution.[12] Because the approximation is quite poor under discretion, we assume that agents' beliefs nest the true rational expectations equilibrium under discretion, so that a set of time-varying "constants" replace x_L and π_L in agents' forecasting model (6). The appendix provides details.

3.3 Benchmark Results

The three panels in Figure 2 plot dynamics for inflation, output and interest rates in response to natural rate shock. In each panel results are shown for rational expectations in black and learning in blue. For each belief assumption, there are 60 possible trajectories. For clarity,

[11] To facilitate comparison with earlier results in the literature, our calibration follows Eggertsson and Woodford (2003).

[12] While Eggertsson and Woodford (2003) choose a truncation point of 60 quarters when studying optimal commitment policies under rational expectations, when comparing to dynamics under discretion, they plot the discretion solution assuming no truncation — there is a constant probability of reverting to steady state into the infinite future and equilibrium outcomes are constant in the low state. This assumption is quantitatively important: it over states the size of the contraction under optimal discretion.

Figure 2: Optimal Policy under Rational Expectations and Learning

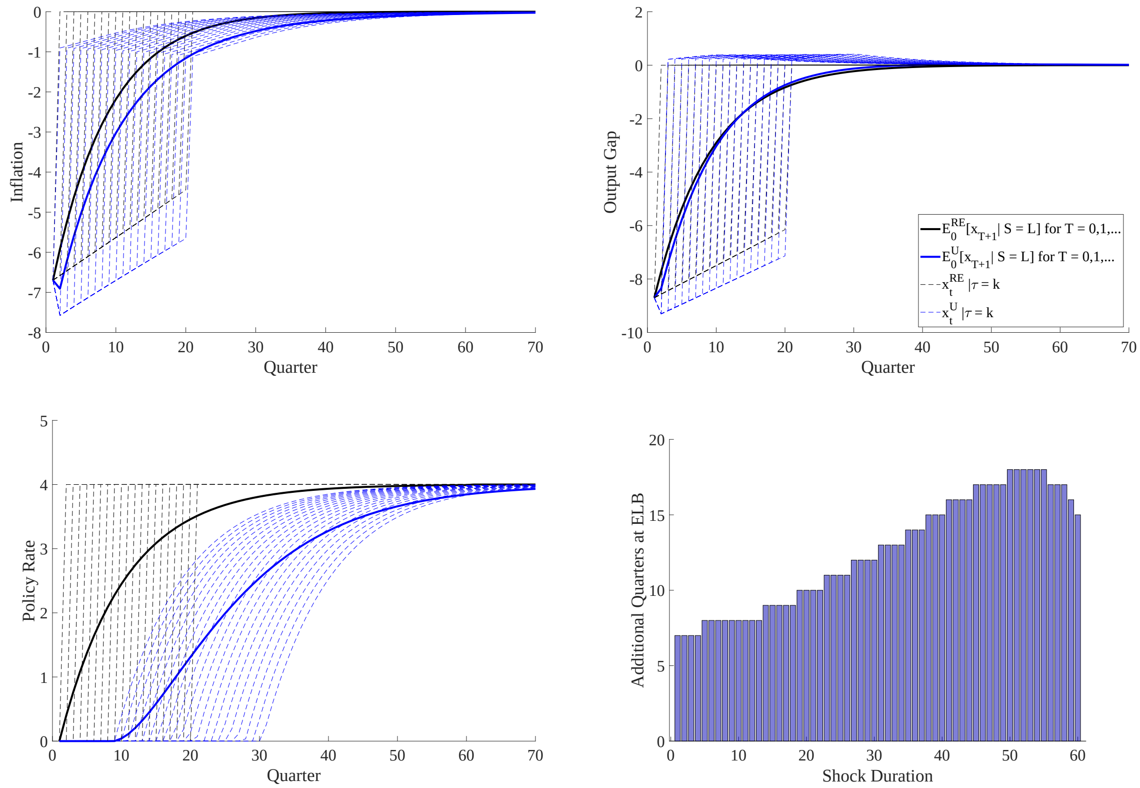

Notes: Parameter values $g = 0.025$, $\rho = 0.9$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$,λ_x, $r_L = -0.005$.

we show the evolution of the economy for the first 20 possible realizations of uncertainty. The solid black and blue lines provide the expected path of each variable conditional on the shock, computed using true data-generating process. These paths therefore summarize the dynamics that the central bank expects at the time of the shock.[13] They are plotted throughout the paper to give a summary of the average characteristics of certain variables given the assumed uncertainty. The final panel plots a histogram of the additional periods at the zero lower bound under learning relative to rational expectations.

Start with optimal discretion under rational expectations. The possible paths of the economy are given by the black lines. When the shock returns to steady state, the economy returns to steady state: monetary policy raises interest rates when it is feasible, and output and inflation are completely stabilized. For realizations of the shock with longer duration, the size of economic contraction declines over time. This is due to the perfect foresight assumption that the natural rate returns to steady state after 60 quarters. The closer the economy to this exit point, the smaller the economic contraction. Agents anticipate a return to normal times, with concomitant income and substitution effects providing increasing sup-

[13] Formally, these are calculated using analogues to equation (8).

port to aggregate demand. The sequence of realizations for output and inflation along these black trajectories are the time-varying constants appearing in agent's beliefs under learning.

Now compare the dynamics under optimal policy with learning. Under learning there are two components of beliefs: the time-varying constants and the estimated drifts. The first component is exogenous and not affected by policy. Monetary policy steers the economy to the extent that it affects the second component, the estimated drifts.

Three insights emerge. First, the impact effect is identical under rational expectations and learning. This reflects the assumption that beliefs are initially consistent with the rational expectations equilibrium in the high state. Consistent with Eusepi and Preston (2011), in subsequent periods extrapolation bias amplifies and propagates the effects of the negative demand shock relative to rational expectations. This is evident in the further fall in aggregate demand and inflation in period two. Second, once the natural rate shock reverts, the optimal policy under learning generates a boom in output. This boom serves to raise inflation and inflation expectations easing the effects of downward drift in inflation expectations on anticipated real interest rates which are too high. In this way the optimal policy under learning shares features of the optimal commitment policy under rational expectations discussed later.[14] Third, under learning the economy spends additional time at the zero lower bound. Even when the natural rate returns to its steady state value, the central bank must maintain zero interest rate policy because falling inflation expectations lead to excessively high real interest rate expectations which constrain demand.

The final panel makes this clear. As the realization of the demand shock becomes more persistent the economy experiences extended periods of zero interest rate policy. This effect moderates for very long shock durations because of the rational component of beliefs. Figure 2 explains why, showing inflation and output gap expectations in the top panels and the associated forecast errors in the bottom panels. For visual clarity we plot realizations occurring four quarters apart. Focus on inflation, understanding that similar comments apply to the output gap. The first panel reveals a critical property of the model. As the duration of the shock lengthens inflation expectations rise monotonically. This reflects the rational component of expectations. However, for longer duration shocks, at the time natural rate returns to the high state, inflation expectations become progressively more pessimistic in that high state. This reflects the endogenous drift component of expectations. Because of the ZLB the decline in inflation expectations triggered by the negative natural rate become self-confirming, producing the observed downward trend. The trend is reversed when the shock ends and the zero interest rate policy produces stimulus, driving up expectations and

[14] As shown in Eusepi, Giannoni, and Preston (2019) optimal policy displays this feature even when the monetary authority is not constrained by the zero lower bound.

realized output and inflation.

The forecast errors in the bottom panel tell the same story. After the initial large surprise, inflation forecast errors become progressively larger over time, though the effect is modest when compared to the initial surprise. The forecast errors lead to downward revisions in beliefs about inflation which cause the zero lower bound to be a constraint even after the shock dissipates. The larger the downward drift in expectations the more binding is the constraint of beliefs on monetary policy, increasing the time at the zero lower bound. However, time at the zero lower bound period moderates for shock durations approaching the truncation point. Here the rational component of beliefs, a stabilizing force driving the economy back to the high state rational expectations equilibrium, dominates the contractionary effect from learning. Forecast errors decline and the downward revision to inflation beliefs slow. This explains the non-monotonicity of the final panel of figure 1.

Figure 3: Beliefs and Forecast Errors under Optimal Policy

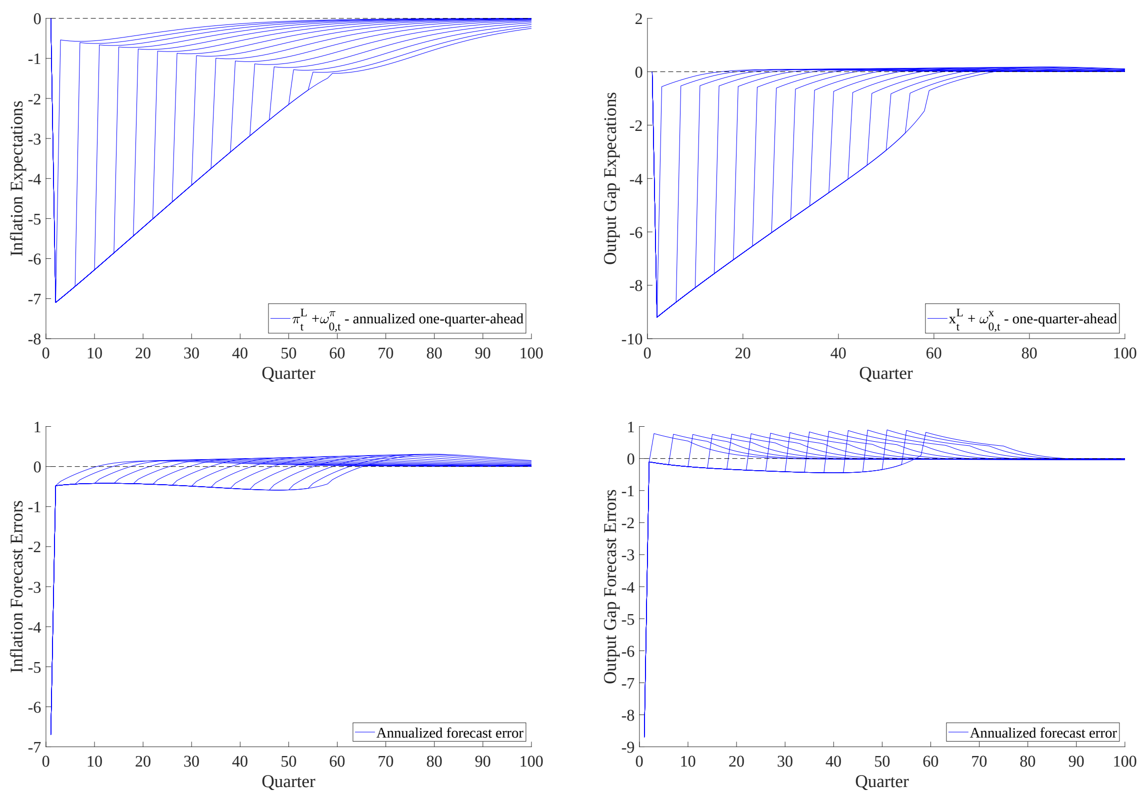

Notes: Parameter values $g = 0.075$, $\rho = 0.985$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\lambda_x = 0.06$, $\kappa = 0.024$, $r_L = -0.005$.

3.4 The Effects of Expectations Anchoring

An important practical concern of central banks since the global financial crisis regards possible instability in long-term inflation expectations. If expectations become unanchored from the inflation target and drift downwards during periods of declining economic activity, the zero lower bound constraint becomes more severe.

The degree to which expectations are well anchored is potentially relevant to understanding the different experiences of Japan and the United States when confronted by the constraint of the zero lower bound. Japan's recession in the early 1990s occurred at a time when inflation expectations in industrialized countries were poorly anchored. In contrast, the global financial crisis occurred at time when most advanced economies had credible inflation targets. As made clear below, our modeling framework provides an explanation of Japan's poorer economic performance.

A natural measure of how well anchored are inflation expectations is the sensitivity of medium to long-term inflation expectations to short-run forecast errors. That is, the Kalman gain. If there is a high degree of sensitivity, so that a negative inflation forecast error leads to a large markdown in long-term inflation expectations, then expectations are poorly anchored. If there is a low degree of sensitivity, inflation expectations are well anchored. Carvalho, Eusepi, Moench, and Preston (2019) and Eusepi, Giannoni, and Preston (2019) formally study the consequences of expectations anchoring for inflation dynamics and monetary policy in normal times. Together these papers adduce evidence that inflation expectations are well anchored in the United States after the late 1990s and that this greatly improves the ability of monetary policy to stabilize the economy. At the same time they show that long-term expectations in Japan have remained poorly anchored since the mid-1990, while some European countries have experienced episodes of de-anchoring at the outset of the 2009 financial crisis.

Given the importance of the expectations formation mechanism to our central results we explore the effects of varying the gain which regulates the sensitivity of long-term inflation expectations to short-run forecast errors. We can also interpret variations in the *persistence* of beliefs, measured by the parameter ρ, as reflecting how well anchored are inflation expectations. For a given sensitivity of inflation expectations to short-run forecast errors, the greater is the persistence in beliefs the longer will long-term inflation expectations be away from target.

Figure 4 shows the evolution of inflation drifts and expectations in response to the natural rate shock. The top row gives the expected evolution of the endogenous drift at the time of the shock. The second row plots the realized inflation drift in the case of a 60 quarter duration shock. The final row shows the additional time at the lower bound for each realization of uncertainty. The left column shows the effects of varying the persistence of beliefs for a fixed Kalman gain equal to our benchmark calibration. The right column fixes the subjective persistence of the drifts to be highly persistent and varies the gain.

Figure 4: Unanchored Expectations - Varying ρ and g

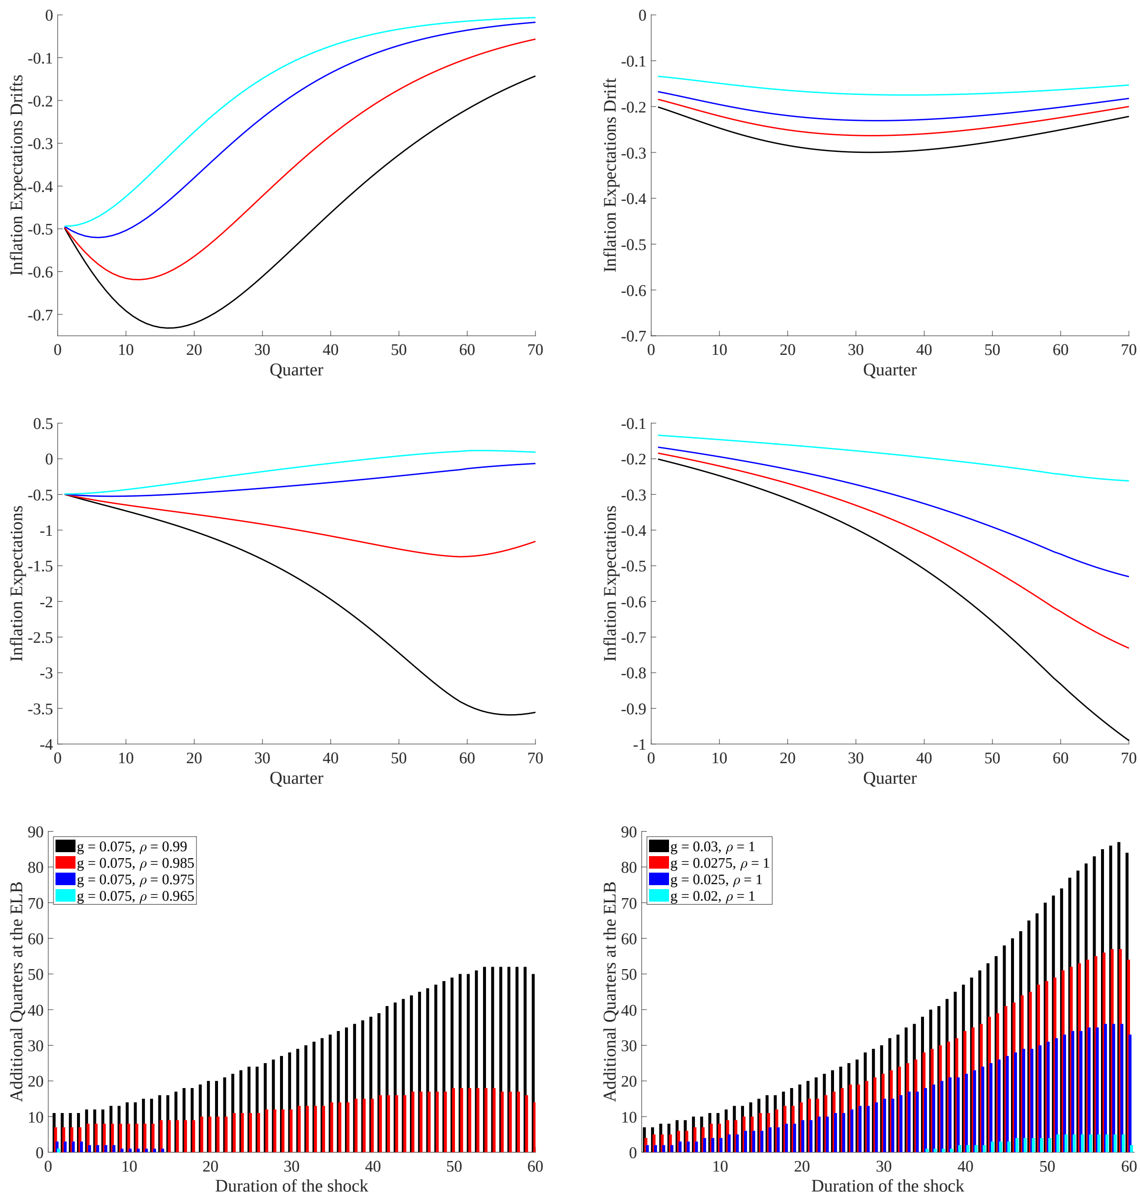

Notes: Parameter values: $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$,λ_x, $r_L = -0.005$.

Starting with the left column, for a fixed sensitivity of beliefs to forecast errors, varying the persistence of subjective beliefs has large effects on equilibrium outcomes. In the second panel, the persistence of beliefs slows the return of inflation expectations to steady state.

And the more sluggish the return in beliefs the bigger effects on dynamics: recall the interest rate elasticity of demand depends critically on this persistence parameter—monetary policy is less effective when beliefs are more persistent, all else equal. Consistent with this, the first panel reveals the progressive fall in inflation expectations. Note that for any realization of uncertainty, beliefs follow this trajectory until the natural rate reverts. It is clear that less well anchored are expectations, the greater the fall in the estimated drift across all realizations of uncertainty.

The final panel underscores the economic consequences: poorly anchored expectations lead to greater the time at the zero lower bound. But modest falls in persistence leads to dramatic falls in the time at the zero lower bound. This not surprising given that, for example, a value of $\rho = 0.95$ implies a half-life in the drift process just over of three years, while a value of $\rho = 0.99$ results in a half-live of nearly twenty years. The latter calibration then induces shifts in expectations at significantly longer horizons; this shows that the degree of anchoring of *long-term expectations* is a crucial determinant in a policy response at the ZLB.

Turning to the right column, with unit root subjective beliefs, equilibrium inflation expectations are highly persistent. This persistence interacts with sensitivity of beliefs to forecast errors with dramatic results. Larger but modest-sized gains lead to prolonged periods at the zero lower bound. Alternatively stated, higher gains represent expectations that are progressively less well anchored. The bottom panel of the column gives a sense of the implications for aggregate dynamics and interest rate policy. As expectations become increasingly unstable, the additional time at the zero lower bound becomes larger. Importantly, this additional time at the zero lower bound is not a deflation trap as is commonly found in papers with adaptive learning. Dynamics in these periods are similar to those shown in Figure 2 with output and inflation persistently, but not significantly, away from steady state.

4 Simple Analytics of Forward Guidance

We have shown that unanchored expectations represent a serious constraint on conventional interest rate policy. The remainder of the paper explores whether better equilibrium outcomes can be achieved by a credible commitment to future zero interest rate policy. This section provides some simple analytics of forward guidance policy. An important result is that our model does not display a forward guidance puzzle. Furthermore, relative to other models of bounded rationality that solve the forward guidance puzzle, learning generates new implications for the conduct of policy. Both the *size* and *timing* of the promised zero interest rate policy are crucial elements of the optimal plan.

Modeling forward guidance. Consider the following thought experiment. The natural rate of interest is constant at its steady-state value for all time, so that $r_t^n = \bar{r}$ for all t. The central bank announces a monetary policy in which nominal interest rates are reduced to zero. Call this the low state, L. With probability $0 < 1 - \nu < 1$ interest rates remain at zero each period. With probability ν they revert to steady state, $R = \bar{r}$. Call this the high state, H. The policy is perfectly credible and understood by agents. The expected duration of zero interest rate policy is then ν^{-1}. Varying ν then allows study of the consequences of different forward guidance policy.

Rational expectations. Under rational expectations agents fully understand the state-contingent announcements about the interest rate. They perfectly observe the states and the probabilities that describe the two-state-Markov process. The central bank can credibly communicate the policy, and fulfills these commitments. Under rational expectations output and interest rate beliefs satisfy

$$\begin{aligned} \hat{E}_t\left(x_{t+1}|S = L\right) &= (1-\nu)x_L + \nu x_H \\ \hat{E}_t\left(\pi_{t+1}|S = L\right) &= (1-\nu)\pi_L + \nu \pi_H \end{aligned}$$

and interest rate expectations satisfy

$$\hat{E}_t\left(R_{t+1}|S = L\right) = (1-\nu) \times 0 + \nu\bar{r}.$$

Standard calculations show there is a rational expectations equilibrium with $x_H = \pi_H = 0$ and $R_H = \bar{r}$ in the high state and

$$\begin{aligned} x_L &= \frac{\sigma(\beta(\nu-1)+1)}{\beta(\nu-1)\nu + (\nu-1)\kappa\sigma + \nu}\bar{r} \\ \pi_L &= \frac{\kappa\sigma}{\beta(\nu-1)\nu + (\nu-1)\kappa\sigma + \nu}\bar{r} \end{aligned}$$

and $R_L = 0$ in the low state.[15] Not surprisingly the zero interest rate policy generates a boom. The rational expectations equilibrium exhibits a version of the forward guidance puzzle. As the expected duration of zero interest rate policy rises without bound, so too does the increase in output.

Result 1. *In the neighborhood of fixed prices, when $\kappa \to 0$, the rational expectations solution is*

$$x_L = \frac{\sigma}{\nu}\bar{r}$$

As $\nu \to 0$ then $x_L \to \infty$.

[15] Again a unique bounded rational expectations equilibrium requires denominator to be positive

The limit result doesn't depend on being the neighborhood of a fixed price equilibrium. The assumption serves to simplify expressions for the output gap as a function of the average duration of forward guidance policy. This facilitates comparison across models.

Imperfect information. As under rational expectations, agents fully understand the state-contingent announcements about the interest rate. They perfectly observe the states and the probabilities that describe the two-state-Markov process. The central bank can credibly communicate the policy, and fulfills these commitments. However, in contrast to rational expectations, agents do not know how forward guidance is going to affect the equilibrium values of output and inflation at the zero lower bound or their equilibrium values when the economy reverts back to the normal state. This approach to modeling central bank communication in adaptive learning models was first adopted by Preston (2006) and Eusepi and Preston (2010).

To clarify, recall the consumption decision rule (1):

$$c_t(i) = \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) y_T - \beta \left(R_T - \pi_{T+1} - r_T^n \right) \right].$$

Consider the case where $\omega^{\pi}_{t|t-1} = \omega^{x}_{t|t-1} = 0$. To evaluate the impact effect of a forward guidance announcement, expectations for output and inflation are fixed at the rational expectations equilibrium *before* the announcement so that, after imposing goods market clearing, the effect of zero interest are policy is given by

$$\begin{aligned} y_t &= -\sigma \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left(R_T - \bar{r} \right) \\ &= \frac{\sigma}{1-\beta\left(1-\nu\right)} \bar{r}. \end{aligned}$$

Because agents don't contemplate how changing future interest rates affect inflation and output gap forecasts, this quantity corresponds to the impact effect of forward guidance in the case of level-1 reasoning—see Farhi and Werning (2019) and García-Schmidt and Woodford (2019). Researchers have used the level-k model of bounded rationality to study how expectations react to a change in policy. Starting from an initial condition, here the rational expectations equilibrium in the high state, agents form expectations about changes in future macroeconomic variables based on a finite deductive procedure about others' behavior. In each successive iteration, agents form expectations using equilibrium outcomes from the previous iteration. When the deductive procedure involves k iterations we have level k reasoning. This process converges to rational expectations as k grows. Here we stop at the first iteration, $k = 1$.[16]

Reintroducing the estimated drifts generates a key difference to existing literature. For simplicity, but without loss of generality, focus on the case where agents estimate a single drift, $\omega^{\pi}_{t|t-1}$, over time. Output gap expectations remain constant, equal to zero. Then inflation expectations satisfy

$$\hat{E}_t\left(\pi_{t+1}|S = L\right) \;\; = \;\; (1-\nu)\omega^{\pi}_{t|t-1} + \nu\omega^{\pi}_{t|t-1} = \omega^{\pi}_{t|t-1}.$$

Expectations about the nominal interest rate satisfy[17]

$$\hat{E}_t\left(R_{t+1}|S = L\right) \;\; = \;\; (1-\nu)\times 0 + \nu(\bar{r} + \omega^{\pi}_{t|t-1}).$$

As before, agents know the state, and understand interest rates are equal to zero in the low state. Now the impact effect on output and inflation depends on initial beliefs:

$$\begin{aligned} x^L_t &= \frac{\sigma}{1-\beta(1-\nu)}(\bar{r} + \omega^{\pi}_{t|t-1}) \\ \pi^L_t &= \kappa\left(\frac{\sigma}{1-\beta(1-\nu)}(\bar{r} + \omega^{\pi}_{t|t-1})\right) + \frac{(1-\alpha)\beta}{1-\alpha\beta}\omega^{\pi}_{t|t-1} \end{aligned}$$

where these expressions assume $\rho = 1$. It is immediate that the learning model does not display a forward guidance puzzle.

Result 2*. The imperfect information solution is*

$$x^L_t = \frac{\sigma}{1-\beta(1-\nu)}(\bar{r} + \omega^{\pi}_{t|t-1})$$

As $\nu \to 0$ then

$$x^L_t = \frac{\sigma}{1-\beta}(\bar{r} + \omega^{\pi}_{t|t-1}).$$

Even when the expected duration of zero interest rate policy rises without bound, output dynamics remain bounded. The effective stimulus depends on the steady-state natural rate. The larger this quantity the larger the effective stimulus of zero interest rate policy. But it also depends on inflation beliefs. Falling inflation expectations reduce the effective stimulus from a commitment to zero interest rate policy. This is because agents anticipate higher real interest rates. It is the presence of this term that distinguishes our analysis from all

[16] The central difference between Farhi and Werning (2019) and us is that beliefs respond to policy over time, generating dynamic general equilibrium effects.

[17] Later analysis assumes that the normalizing constants in beliefs are unaffected by forward guidance policy. We therefore set them here to zero for simplicity.

other recent models that resolve the forward guidance puzzle. Policy must prevent inflation expectations falling. Failure to do so means that monetary policy becomes less effective over time.

Of course, the effects of policy can still be large under standard parameterizations of the household's discount factor, β. Indeed, in the language of Farhi and Werning (2019) the factor

$$\frac{\sigma}{1-\beta}$$

is a pure partial equilibrium effect.[18] It captures the direct effect of a shift in real interest rates expectations by the amount

$$\bar{r} + \omega^{\pi}_{t|t-1}.$$

Because the effects of this shift are permanent in this simple thought experiment, the effects of forward guidance policy can still be large. What the sequel demonstrates is that general equilibrium effects that operate through the revision of beliefs considerably complicate the ability of forward guidance policy to deliver effective stimulus.

Empirical evidence. A growing literature supports the two key assumptions underpinning our framework: low-level reasoning and learning from past outcomes. First, agents generally use a low level of deduction, ranging from zero to three across different experiments. Nagel (1995) and Bosch-Domènech, Montalvo, Nagel, and Satorra (2002) analyze beauty contests in lab experiments and surveys and find most people play level one or two. Camerer, Ho, and Chong (2004), Arad and Rubinstein (2012) reach similar conclusions in the context of different types of games played in the lab. Second, Duffy and Nagel (1997) and Ho, Camerer, and Weigelt (1998) study dynamic beauty context games in the lab. In addition to low levels of deduction they show evidence of learning over time, in response to past outcomes. Closer to our framework, Evans, Gibbs, and McGough (2019) find strong evidence supporting our two model assumptions using a lab experiment of an announced structural change.

Forward guidance puzzle. Our paper is not the first to provide a resolution of the forward guidance puzzle. There have been three prominent fixes. The first class of model introduces finite lives and life-cycle considerations (Del Negro, Giannoni, and Patterson 2012 and Eggertsson, Mehrotra, and Robbins 2019). The second class of model introduces bounded rationality and imperfect information (Angeletos and Lian 2018, Woodford 2018 and Gabaix 2020). The third class of models introduce incomplete markets and heterogeneous agents (Bilbiie 2018 and McKay, Nakamura, and Steinsson 2017). Some papers combine elements of each, such as Farhi and Werning (2019) which includes both bounded

[18] As they note, there is a small general equilibrium effect from contemporaneous goods market clearing.

rationality and incomplete markets.

A common property of these models is that they discount the future more heavily to limit the effect of shifting beliefs about future economic conditions on current equilibrium outcomes. To be concrete, in general these papers have the approximate representation of our simple New Keynesian model

$$\begin{aligned} x_t &= ME_t x_{t+1} - \sigma N(R_t - E_t\pi_{t+1} - r_t^n) \\ \pi &= M_f \beta E_t \pi_{t+1} + \kappa x_t \end{aligned}$$

where $0 < M \leq 1$, $0 < M_f \leq 1$ and $0 < N \leq 1$. Full information rational expectations is the special case $M = M_f = N = 1$. The constants M and N are implications of assumptions on incomplete markets, borrowing constraints and fiscal transfers. The constants M and M_f are influenced by myopia and limited planning horizons, level-k reasoning and diverse private information. Regardless of underlying assumptions, the effects are to increase discounting of future equilibrium outcomes and reduce the sensitivity of demand to variations in future real interest rates. Both serve to reduce the impact of forward guidance policy.

Considering the same thought experiment as above provides the following result.

Result 3. *In the neighborhood of fixed prices, when $\kappa \to 0$, the rational expectations solution is*

$$x_L = \frac{\sigma N}{1 - M(1-\nu)}\bar{r}$$

As $\nu \to 0$ then

$$x_L \to \frac{\sigma N}{1-M}\bar{r}$$

The solution for output has the same form as our learning model which is recovered under the parameterizations $N = 1$ and $M = \beta$. The efficacy of monetary policy again depends on the steady-state natural interest rate. And as before this quantity is the partial equilibrium effect in this broad class of nested models. What we later show is that the general equilibrium effects in these nested models remain large, even when the partial equilibrium effects are substantially moderated. An implication is that the optimal forward guidance policy is almost indistinguishable from the rational expectations commitment policy, even though these models do not display a forward guidance puzzle. The predictions of our model are qualitatively and quantitatively different, because of distinct general equilibrium effects.

5 Forward Guidance Policy

Now suppose the central bank can commit to zero interest rate policy in the future. Let τ denote the date at which the natural rate returns to the high state. For each τ the central bank makes a promise of k_τ periods of zero interest rate policy. A forward guidance policy is then the set of promises $\{k_\tau\}$ for $\tau \in [1, ..., 60]$. These state-contingent promises are assumed to be fully credible. After the forward guidance period, monetary policy returns to the optimal policy problem of Section 3.

Consistent with the simple framework presented above, we assume that forward guidance leads households to revise only their beliefs about future interest rate policy. Expectations about future inflation and the output gap remain unchanged in response to the announced forward guidance policy. However, these beliefs will be revised in future periods in response to changed macroeconomic conditions that result from the announced policy. This assumption reflects evidence from surveys of professional forecasters' expectations in Crump, Eusepi, and Moench (2011), Del Negro, Giannoni, and Patterson (2012), Campbell, Evans, Fisher, and Justiniano (2012) and Andrade and Ferroni (2021). For example, Crump, Eusepi, and Moench (2011) show that in response to the Federal Reserve's changed forward guidance in 2011, the cross-sectional average term structure of expectations about future interest rates shifts to being consistent with the announced path for short-term interest rates. At the same time, the distribution of expectations across forecasters compresses substantially around the announced path. In contrast, there are only modest changes to the average term structure of expectations for inflation and output and their distribution across individual forecasters.

5.1 The policy problem

The central bank chooses forward guidance policy to minimize the loss function (17) subject to the constraints implied by household and firm behavior. The appendix writes down the optimal policy problem, first-order conditions and describes the solution algorithm. The solution has the following general characteristics. As before, there are three regimes. The first is defined by the natural rate at r_L and interest rates at the ZLB. The central bank makes a set of state-contingent credible promises. The second is defined by the natural rate reverting to r_H but interest rates remaining at zero, consistent with the announced commitment. The third is defined by the economy being in the high state and the central bank using conventional interest rate policy.

To give some feel for the economics of the solution, consider the structural equations for aggregate demand and supply across regimes. The third regime occurs in periods $t > \tau + k_\tau$ where τ is the duration of the shock and k_τ the period of zero interest rate policy attached to that state-contingent realization. The dynamics then coincide with regime 3 from the no forward guidance case and are written

$$x_t \;\; = \;\; -\sigma\left(R_t - r_H\right) + \frac{1-\beta}{1-\beta\rho}\omega^x_{t|t-1} + \sigma\frac{1-\beta}{1-\beta\rho}\omega^\pi_{t|t-1}.$$

In all regimes, forward guidance policy only affects the structure of the aggregate demand relationship. The equation describing aggregate supply is always given by (12).

The second regime occurs during $\tau \leq t \leq \tau + k_\tau$, when the natural rate is r_H but interest rates are zero — the zero interest rate policy continues beyond the shock. This is the second regime. The appendix shows that aggregate demand satisfies

$$\begin{aligned} x_t^{k_\tau(j)} &= -\sigma\left[\beta\left(\frac{\beta^{k_\tau-j}}{1-\beta}r_H + \frac{(\beta\rho)^{k_\tau-j}}{1-\beta\rho}\omega^\pi_{t|t-1}\right) - \frac{\sigma}{1-\beta\rho}\omega^\pi_{t|t-1} - \frac{1}{1-\beta}r_H\right] + \frac{1-\beta}{1-\beta\rho}\omega^x_{t|t-1} \\ &= \sigma\left[\frac{1}{1-\beta} - \beta^{k_\tau-j}\frac{\beta}{1-\beta}\right]r_H + \frac{1-\beta}{1-\beta\rho}\omega^x_{t|t-1} + \left(1-\beta\left(\beta\rho\right)^{k_\tau-j}\right)\frac{\sigma}{1-\beta\rho}\omega^\pi_{t|t-1}. \end{aligned}$$

After the first equality, the term in brackets measures the difference between the present discounted value of the real rate and the natural rate of interest, while the second term captures expected income. After the second equality, we see there are two sources of stimulus from the zero interest rate policy. The first term captures the reduction in nominal rates to zero relative to the real neutral rate for $k_\tau - j$ periods. Recall the standard one period effect is σr_H. The first term is simply the sum of this effect over $k_\tau - j$ periods. Of course, what matters is the effective real interest rate relative to the neutral rate. This is determined by accounting for inflation expectations in the final term: the first component is positive, as lower inflation expectations increase the expected real rate, while the second term is negative. This latter term, from the Fisher equation, captures the expansionary effect from returning to a lower expected long-term nominal interest rate. On net, lower inflation expectations lower aggregate demand. Finally, the stimulus from zero interest rate policy declines as the economy approaches liftoff.

The first regime occurs during $t < \tau$ when the natural rate is in the low state. This is the most complicated regime because to determine aggregate demand we must account for the state-contingent character of forward guidance policy to be implemented in all future realizations of uncertainty. To compute aggregate demand in the low state we work recursively backwards in time. At $t = \tau = \bar{T}$ the economy returns to the high state with probability one. At that time the central bank implements $k_{\bar{T}}$ periods of zero interest policy.[19] Using this fact and that interest rate forecasts must satisfy

$$E_t R_{t+1} = r_H + \omega^\pi_{t|t-1}$$

[19]Notice this is the first period of regime 2 so it corresponds to $k_{\bar{T}}(0)$, that is when $j = 0$.

for $t \geq \bar{T} + k_{\bar{T}}$ permits

$$E_{\bar{T}} \sum_{T=\bar{T}}^{\infty} \beta^{T-t} R_{T+1} = \frac{\beta^{k_{\bar{T}}}}{1-\beta} r_H + \frac{(\beta\rho)^{k_{\bar{T}}}}{1-\beta\rho} \omega^{\pi}_{\bar{T}|\bar{T}-1}.$$

Working progressively backwards in time, the appendix shows this expression generalizes to

$$E_{\bar{T}-j} \sum_{T=\bar{T}-j}^{\infty} \beta^{T-\bar{T}-j} R_{T+1} = \Psi^{\bar{T}-j}_{\rho,k_{\bar{T}-j}} \omega^{\pi}_{\bar{T}-j|\bar{T}-j-1} + \Psi^{\bar{T}-j}_{k_{\bar{T}-j}} r^H$$

where

$$\begin{aligned}
\Psi^{\bar{T}-j}_{\rho,k_{\bar{T}-j}} &= (1-\delta)\,\beta\Psi^{\bar{T}-j+1}_{\rho,k_{\bar{T}-j+1}} + \delta\frac{(\beta\rho)^{k_{\bar{T}-j}}}{1-\beta\rho} \\
\Psi^{\bar{T}-j}_{k_{\bar{T}-j}} &= (1-\delta)\,\beta\Psi^{\bar{T}-j+1}_{k_{\bar{T}-j+1}} + \delta\frac{\beta^{k_{\bar{T}-j}}}{1-\beta}
\end{aligned}$$

for $j = 1, ..., \bar{T}$ and

$$\begin{aligned}
\Psi^{\bar{T}}_{\rho,k_{\bar{T}}} &= \frac{(\beta\rho)^{k_{\bar{T}}}}{1-\beta\rho} \\
\Psi^{\bar{T}}_{k_{\bar{T}}} &= \frac{\beta^{k_{\bar{T}}}}{1-\beta}
\end{aligned}$$

for $j = 0$.

These recursions encode the state-contingent properties of forward guidance. For example, when $j = 1$

$$\begin{aligned}
\Psi^{\bar{T}-1}_{\rho,k_{\bar{T}-1}} &= (1-\delta)\,\beta\Psi^{\bar{T}}_{\rho,k_{\bar{T}}} + \delta\frac{(\beta\rho)^{k_{\bar{T}-1}}}{1-\beta\rho} \\
&= (1-\delta)\,\beta\frac{(\beta\rho)^{k_{\bar{T}}}}{1-\beta\rho} + \delta\frac{(\beta\rho)^{k_{\bar{T}-1}}}{1-\beta\rho}.
\end{aligned}$$

There are two possible realizations of uncertainty: remain in the low state for one period then implement forward guidance promise of $k_{\bar{T}}$ or return to the high state in the period $\bar{T}-1$ and implement forward guidance of $k_{\bar{T}-1}$. When j periods before $\bar{T}$ there are 2^j such terms, capturing all possible paths to the high state and the associated forward guidance policy. Note that the larger the promises $\{k_{\bar{T}-1}, k_{\bar{T}}\}$ the smaller is $\Psi^{\bar{T}-1}_{\rho,k_{\bar{T}-1}}$; that for constant promises $\{k, k\}$ the fixed amount of forward guidance is more effective when implemented in period $\bar{T}-1$ rather than $\bar{T}$; and that for a flat profile of promises $\{k, k, ..., k\}$ the term $\Psi^{\bar{T}-j}_{\rho,k_{\bar{T}-j}}$ is rising in j. Together this means that earlier and larger action is more stimulatory.

Aggregate demand then can be written as

$$\begin{aligned} x_t^L &= x_{L,t}^{RE} + \frac{(1-\beta)}{1-\beta\rho}\omega_{t|t-1}^x + \frac{\sigma}{1-\beta\rho}\omega_{t|t-1}^\pi + \sigma\beta\left[\frac{1}{1-\beta} - (1-\delta)\frac{1-[\beta(1-\delta)]^{\bar{T}-t+1}}{1-\beta(1-\delta)}\right] r_H \\ &\quad -\sigma\beta\left(\Psi_{\rho,k_{\bar{T}-t}}^{\bar{T}-t}\omega_{t|t-1}^\pi + \Psi_{k_{\bar{T}-t}}^{\bar{T}-t} r_H\right). \end{aligned}$$

The top line is now familiar. The first term is the normalizing constant given by the rational expectations solution under discretion. This coefficient is not affected by forward guidance. The other terms measure the effects coming from the anticipated changes in output, inflation and the real neutral rate, for a *given* expected path of the nominal interest rate. The lower the expected paths of these variables, the lower aggregate demand.

The bottom line measures the effect from the expected path of the nominal interest rate with state-contingent forward guidance. For the benefit of exposition, assume $\rho = 1$ so we can write the term more compactly as

$$-\sigma\beta\Psi_{k_{\bar{T}-t}}^{\bar{T}-t}\left(\omega_{t|t-1}^\pi + r_H\right).$$

When the 'terminal' nominal interest rate, $\omega_{t|t-1}^\pi + r_H$, is positive, forward guidance should make large promises to make $\Psi_{k_{\bar{T}-j}}^{\bar{T}-j}$ small to maximize the stimulus from zero interest rates. For a given amount of forward guidance, a lower drift (here $\omega_{t|t-1}^\pi < 0$) implies more stimulus. That is because the terminal nominal interest rate is expected to revert back to a lower long-run equilibrium value, implying a flatter path for the nominal interest rate. Again, this effect comes from the Fisher equation.

Of course, the overall effect of the inflation drift on aggregate demand depends on expected real interest rates which are measured by

$$\sigma\left(\frac{1}{1-\rho\beta} - \beta\Psi_{\rho,k_{\bar{T}-t}}^{\bar{T}-t}\right)\omega_{t|t-1}^\pi.$$

Any downward movement in the inflation drift has a negative impact in aggregate. Zero interest rate policy is simply less effective. In response to a drop in the natural rate of interest it is then desirable to announce a forward guidance policy as early as possible, before beliefs start their downward drift. As we show in section 6.3 the longer a central bank waits, the larger the forward guidance commitment is required.

5.2 Results

Figure 5 plots dynamics under the optimal forward guidance policy and learning. As in earlier figures, we plot only the first 20 of 60 possible realizations of the natural rate disturbance. The first three panels plot inflation, the output gap and the policy rate. The black lines show the trajectory of the economy with forward guidance. The blue lines show the trajectories without forward guidance, reproducing earlier results for comparison. The solid blue and black lines give the expected trajectory at time 0 under each policy. The fourth panel shows the additional time at the zero lower bound with and without forward guidance. The final two panels plot inflation and output gap expectations.

A credible commitment to zero interest rate policy dramatically improves equilibrium outcomes. The solid blue and black lines reveal the general character of policy that the central bank implements. Both inflation and output fall by less on average, and recover more quickly even in the face of persistent shocks. Indeed, the optimal forward guidance policy generates a boom in output and an over-shooting of inflation. The over-shooting is subsequently unwound by restraining aggregate demand, delivering a contraction in real activity. Central to this narrative are inflation expectations. Falling inflation expectations make the zero lower bound constraint on nominal interest rates more severe. Optimal policy not only arrests the fall in inflation expectations, but in fact raises inflation expectations relative to steady state. This "reflation" provides additional scope for monetary policy stimulus even after the period of zero interest rate policy, since what matters is the expected real interest rate relative to the neutral rate.

To restrain inflation expectations requires a contraction in real activity. This is implemented by a rise in nominal interest rates after the period of zero interest rate policy. At the maintained parameter values, these interest rate hikes are substantial. However, the important point is there is a fundamental trade-off: stabilizing inflation requires generating a boom in output. This boom prevents inflation expectations from falling, but generates overshooting in inflation which must ultimately be unwound by inducing a contraction. This trade-off underscores the limits of central bank policy under imperfect information. Central bank communication, while perfectly credible, cannot fully shape inflation and output expectations, as in the case of full information. Nonetheless central bank communication affects expectations beyond the horizon of the crisis and forward guidance period (beliefs are state variables). Inflation expectations overshoot across forecasting horizons and the central bank has to respond. We later return to this trade-off, analyzing its dependence on central bank preferences and subjective beliefs.

The fourth panel shows the profile of promises that implements the optimal forward

Figure 5: Unanchored Expectations: Optimal Policy

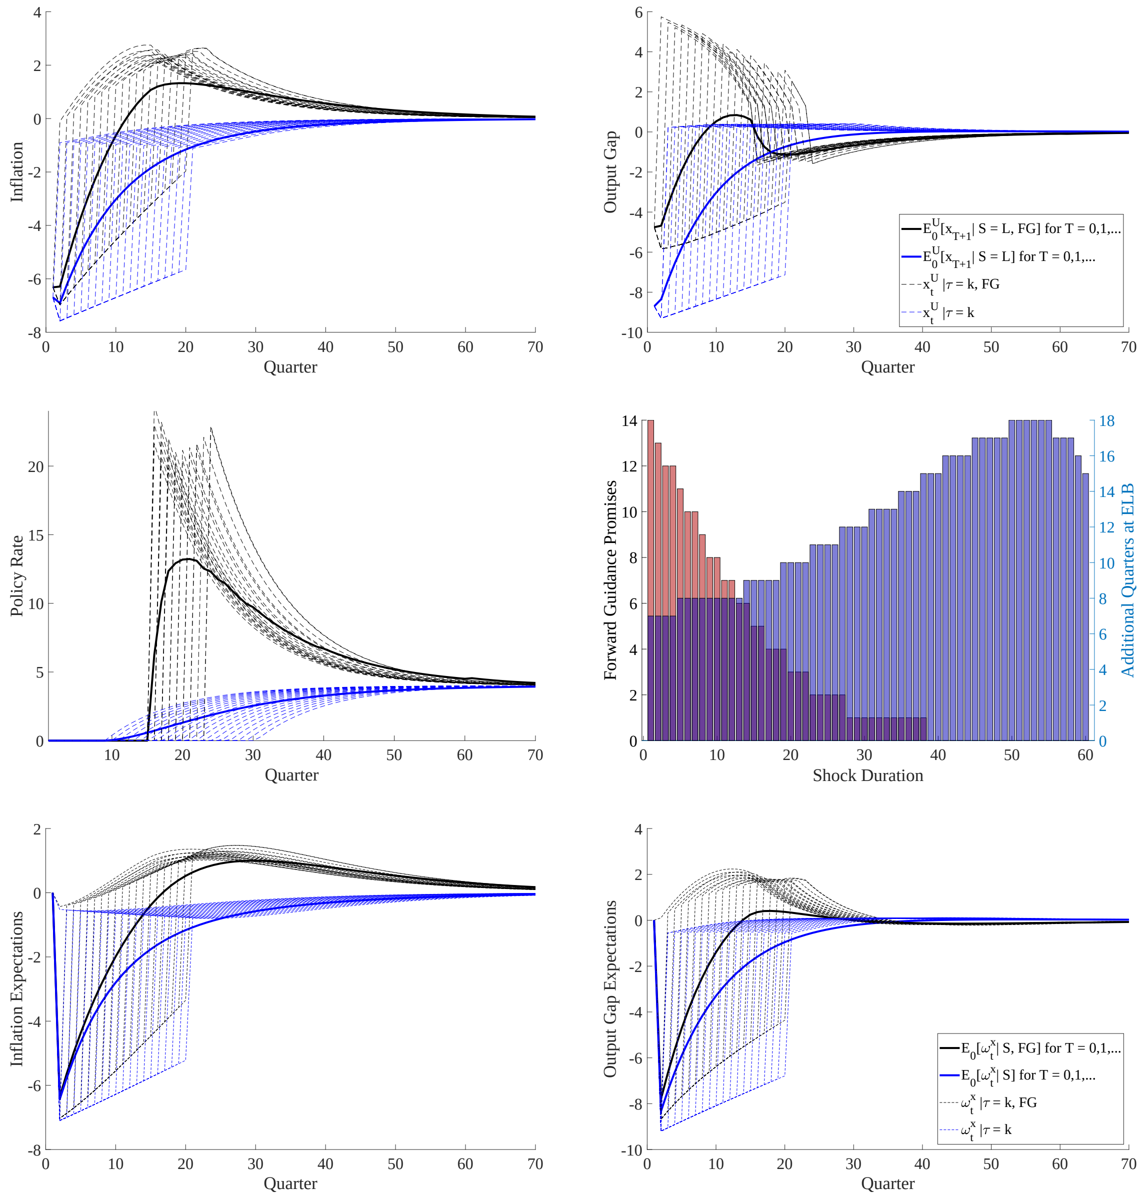

Notes: Parameter values $g = 0.075$, $\rho = 0.985$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.024$,, $\lambda_x = 0.06$, $r_L = -0.005$.

guidance policy. Again, for comparison, we show the additional periods at the zero lower bound under the no forward guidance policy. The profile is large and front loaded. At the time of the shock the central bank commits to substantial stimulus even in the case of short-duration shocks, with the amount of stimulus gradually declining for longer-duration shocks. For shocks that last more than 40 quarters it is optimal to provide no commitment to zero interest rate policy.

Optimal policy therefore displays an insurance principle. The central bank makes large state-contingent promises for short-duration shocks to ensure inflation expectations don't fall in the event of a long-duration shock. The return on these aggressive announcements are the rise in inflation expectations. This affords precious nominal space even after the period of zero interest rate policy. Indeed, this nominal space from higher inflation expectations explains why the optimal forward guidance policy does not promise to keep interest rates at zero for shocks lasting 40 or more quarters. The favorable movement in beliefs removes the constraint of the zero lower bound.

Of course, should the economy experience a favorable short-duration shock, the forward guidance commitment has put substantial stimulus in place which creates a boom. The stimulus is substantial because households anticipate large promises *for all realizations of uncertainty* in the near to medium term. Hence the tension: long-duration shocks have large negative consequences for the economy which must be managed by providing large front-loaded promises. But in the event of a short duration shock, inflation and inflation expectations will be too high, requiring a contraction in aggregate demand. This is the price of insurance, the insurance premium.

The size of the boom and large interest rate responses after the period of zero interest rate policy reflect two model properties: i) interest rate beliefs in the high state; and ii) central bank preferences for output gap and inflation stabilization. On returning to the high state after the forward guidance period we assume agents are uncertain about nominal interest rate policy. The fact that beliefs are formed without knowledge of the policy strategy means the central bank has looser control of the expected path of the nominal interest rate. The policy hike does not raise long-term interest rate expectations as quickly. Rising inflation expectations lead to falling real interest rate expectations, a source of stimulus for aggregate demand. For this reason, the burden of stabilization policy falls on the contemporaneous interest rate. However, should agents have knowledge of the policy strategy, then the central bank could better control real interest rate expectations. For example, if agents understand the central bank to use a Taylor rule so that nominal interest rates rise more than one-for-one with inflation expectations, then real interest expectations would rise with an increase inflation expectations. And anticipation of this fact limits the increase in real activity. Tighter control of future interest rate expectations would reduce reliance on current nominal interest rates. In learning models, central bank communication affords stabilization benefits not just at the zero lower bound — see Eusepi and Preston (2010) for discussion.

As for central bank preferences, Figure 6 shows the consequences of increasing this stabilization weight. Progressively greater weights lead to a smaller over-shooting of inflation, a larger boom in output, and more moderate post forward guidance interest rate response on average. Not surprisingly, the subsequent downturn becomes progressively smaller given greater concern for stabilizing output. What is important is that the profile of forward guidance promises is largely unaffected across stabilization weights — there are only minor differences in the large front-loading of promises.

Figure 6: Unanchored Expectations: Optimal Policy and the Central Bank's Loss Function

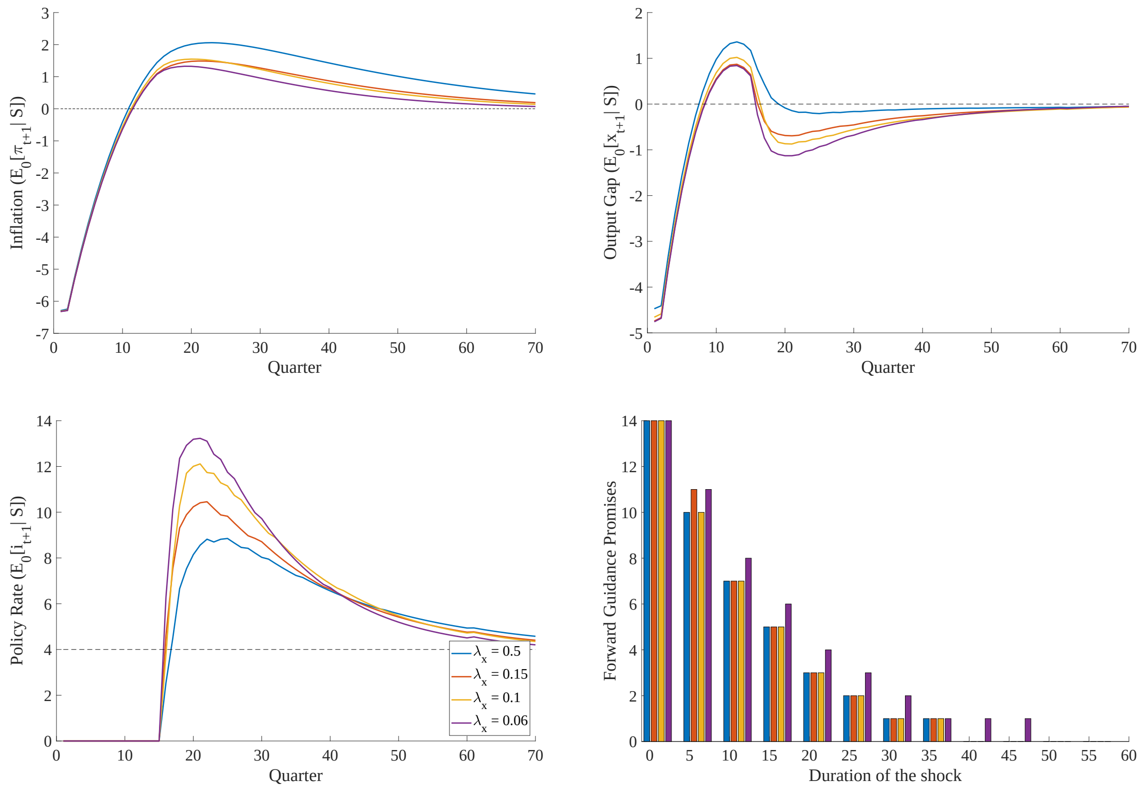

Notes: Parameter values $g = 0.075$, $\rho = 0.985$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.024$, $r_L = -0.005$, $\lambda_x = 0.06$.

6 The Insurance Principle

This final section explores further the economics of the insurance principle. First, we compare the optimal forward guidance policy with unanchored expectations to the optimal policy with rational expectations as well as other recent behavioral New Keynesian models. Despite offering a solution to the forward guidance puzzle, the behavioral models' predictions are almost identical to the rational expectations optimal commitment policy. Second, we shed some light on the nature of the trade-off confronting policy by allowing the central bank to renege on its promises and evaluating the costs of commitment. Third, we explore the costs of delay in implementing the optimal forward guidance policy. This permits insight into the economic outcomes of countries such as Japan that were late to adopt unconventional policy measures at a time when long-term inflation expectations were poorly anchored.

6.1 Alternative models

Using a specific thought experiment, Section 4 compared the properties of our model to number of other papers on zero interest rate policy. We now characterize the optimal forward guidance policy for two of these models and compare them with the optimal commitment policy under rational expectations.

Figure 7: Optimal Policy: RE Commitment, Gabaix's, THANK

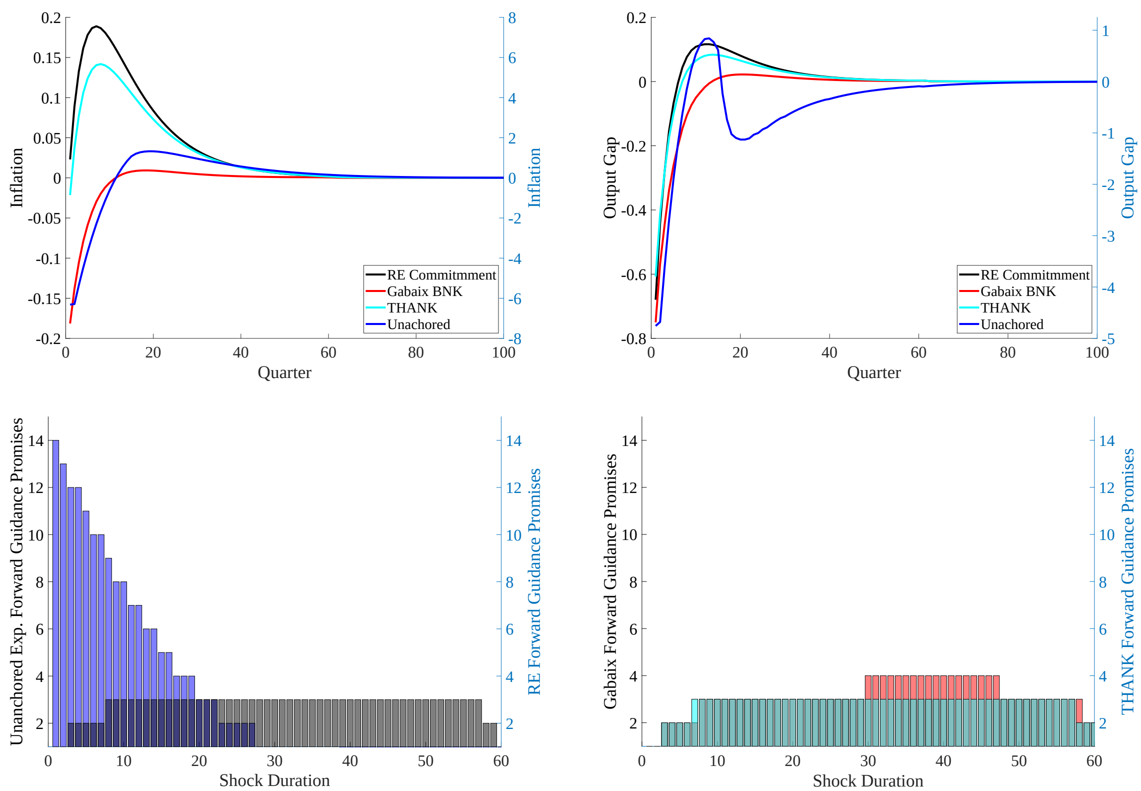

Notes: Shared parameter values $\lambda_x = 0.06$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.024$, $r_L = -0.005$. Gabaix parameters: $M = 0.85$, and $M^f = 0.8$. THANK parameters: $M = 0.9701$, $M^f = 1$, and $N = 0.843$.

Figure 7 summarizes the findings. The top two panels show the dynamics of inflation and output for four different models. As in earlier figures, the plotted variables correspond to the central bank's expectations of these variables at the time of the shock, conditional on optimal policy. The four models correspond to rational expectations (black); Gabaix's (2020) behavioral New Keynesian (BNK) model (red); Bilbiie's (2018) heterogeneous agent (THANK) model (turquoise); and our theory of unanchored expectations (blue). The bottom two panels plot the state-contingent forward guidance promises for each of these models. Since these models introduce some additional parameters, we take values reported to be the preferred specifications by the authors.[20] Because of increased discounting of the future and reduced sensitivity to real interest rates, foward guidance in the THANK and BNK models is less powerful than the standard rational expectations New Keynesian model.

The figure reveals our model is radically different to existing approaches on three grounds. First, as evidenced by the different axes in the figures, optimal policy does not nearly deliver the same stimulus. The same shock delivers a decline in output gap and inflation that is roughly one order of magnitude larger. Second, unanchored expectations induce a substantially different profile of forward guidance promises. While the unanchored expectations models has large front-loaded promises, the other models provide fairly flat profiles, that display a modest hump-shape.[21] Perhaps most striking is how similar the optimal forward guidance policies are across the rational expectations, THANK, and BNK models. Indeed they are almost identical, with minor differences, notably for the BNK model with shock durations between 30 and 45 quarters where an additional one quarter of forward guidance is offered. Third, because of imperfect control of the term structure of expectations, forward guidance delivers a strong rebound driven by long-term expectations. Consequently, it features a tighter monetary policy when exiting the zero lower bound generating a downturn in economic activity.

The relative weakness of forward guidance policy in our framework can be sourced to general equilibrium effects. In our model these general equilibrium effect are muted, as household and firm inflation and output gap expectations do not respond on impact to the forward guidance announcement. Beliefs are revised subsequently in response to observed data which considerably weakens general equilibrium effects. Moreover, poorly anchored expectations produce negative general equilibrium effects that are only gradually counteracted by forward guidance policy. These general equilibrium consequences are the reason for the front-loading of policy stimulus.

These differences cannot be sourced to the forward guidance puzzle. In fact, in Gabaix's model households and firms discount future equilibrium outcomes to a much greater extent than in the rational expectations model; the model is free of the forward guidance puzzle. However, agents still contemplate how forward guidance policy shifts the equilibrium outcomes for output and inflation in those future periods. The resulting general equilibrium income and substitution effects remain large. Indeed, sufficiently large to unwind the smaller

[20] In the case of Bilbiie's (2018), his calibration seeks to match the model of McKay, Nakamura, and Steinsson (2016).

[21] Those familiar with Eggertsson and Woodford (2003) might question the hump-shape profile. In their paper they restrict policies to fall in a class in which promises are monotonically rising with the duration of the shock. We do not impose this restriction. Not surprisingly, the perfect foresight assumption for the end of the natural rate shock then has some consequences for very long duration shocks. Notably, the closer to exit, the smaller the required promises.

partial equilibrium effects from discounting.

Finally, note that our results do not rely on assumptions about equilibrium selection. Cochrane (2017) argues that many features of the NK model are implausible because they rest on equally implausible equilibrium selection criteria. These equilibrium selection concerns are present in the RE, THANK and BNK models, but not in our model. This is because while agents are forward looking, their beliefs are state variables. Formally, this makes our model and old Keynesian model in Cochrane's language.

6.2 The costs of insurance

The insurance principle delivers too much stimulus in the case of a short duration shock. A commitment to fulfill promised monetary accommodation in the event of a favorable shock requires the central bank to restrain economic activity and therefore inflation expectations. To evaluate the costs of such commitment or, in other words, the insurance premium, we consider the following experiment. At the time of the negative demand shock the central bank announces the optimal forward guidance policy. However, when the natural rate shock reverts to steady state, if desirable, the central bank raises interest rates and reneges on the announced zero interest rate policies.[22] The path of the policy rate is then optimally determined, consistent with regime 3 in the earlier commitment problem.

Figure 8 displays the results. The structure of the figure is identical to Figure 5, except for the promises being replaced by output gap beliefs. The blue lines in each panel now show equilibrium outcomes for a central bank that reneges on the announced forward guidance. Early abandonment of zero interest rate policy has significant implications for optimal policy. On average the renege policy generates less over-shooting of inflation and never generates a boom in output. Consistent with this, the bottom panels show the promised stimulus is less entrenched in inflation expectations. Interest rates rise by much less, peaking at around 6 per cent, much below 13 percent for the fully optimal policy. Importantly, the central bank never makes payment on the insurance, avoiding the need to engineer a recession to restrain inflation expectations. Indeed, this policy improves welfare by roughly 20 per cent.

The dashed lines for the individual realizations provide additional nuance. Except for shocks lasting a single period, for short duration shocks the renege policy creates a very small positive output gap at the time the natural rate reverts to steady state. In the case of a one-period shock, it remains optimal to maintain zero interest rate policy — reneging implies a the central bank would choose a negative interest rate. As the duration of the shock rises, the output gap at the time of normalization progressively falls, so that longer duration shocks

[22]It may not be desirable if zero interest rate policy continues to be optimal. In this case, the central bank will raise rates in the first period that it is desirable.

Figure 8: Unanchored Expectations: Optimal Policy vs Optimal Policy Renege

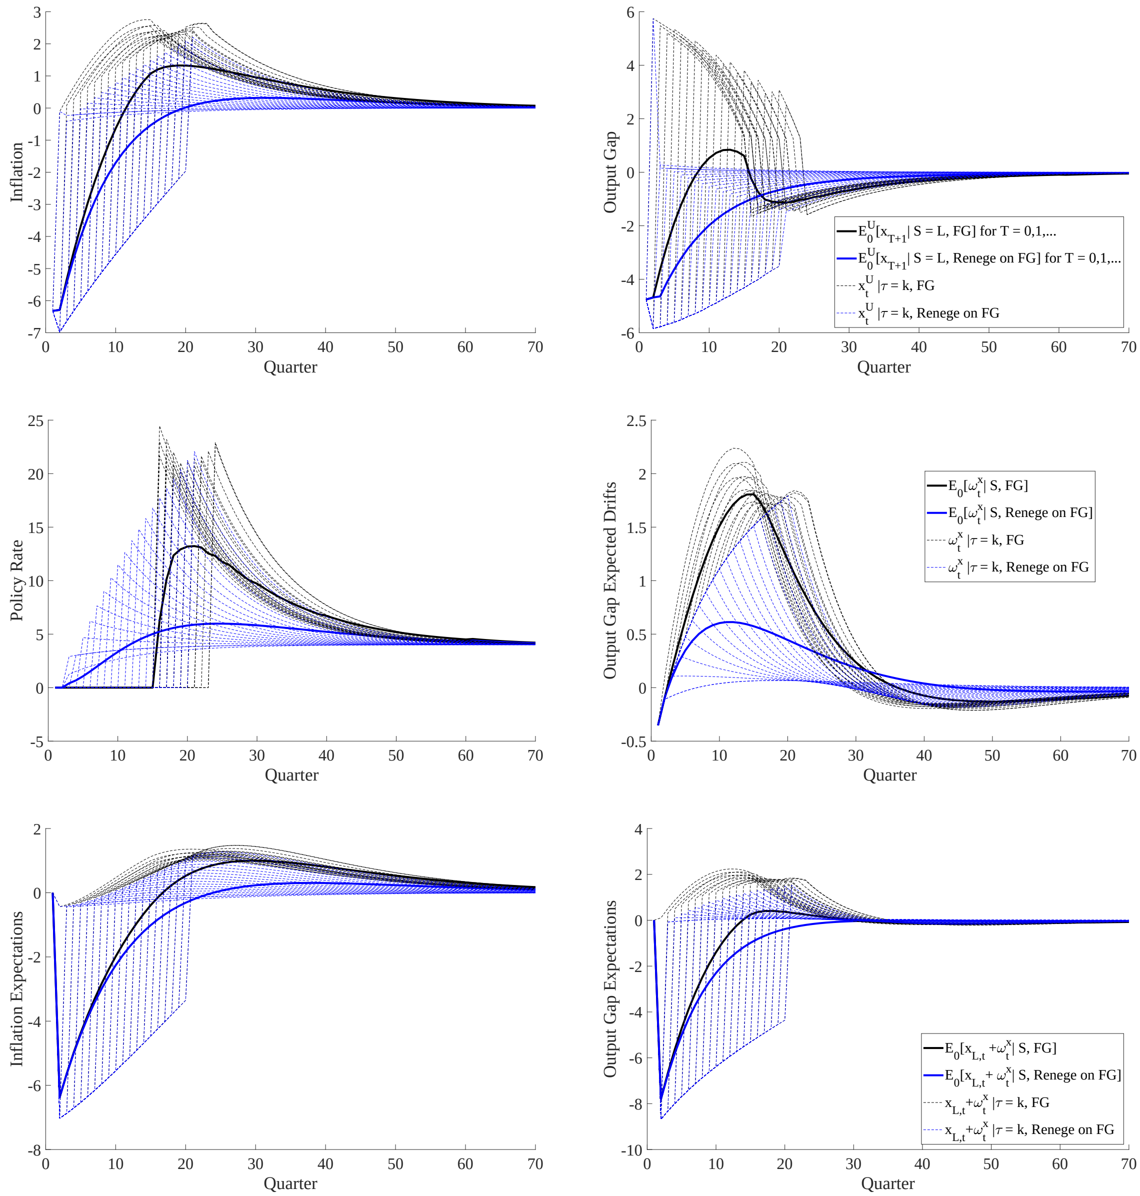

Notes: Parameter values $g = 0.075$, $\rho = 0.985$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.024$,, $\lambda_x = 0.06$, $r_L = -0.005$. There is a $\approx 20\%$ improvement in welfare by reneging on the policy after the shock ends.

never generate a positive output gap. Nonetheless, inflation expectations progressively rise at the time of normalization in monetary policy, with the interest rate response rising also from initially modest increases to more substantial values in the case of medium-duration shocks. With the passage of time, expectations do respond to the anticipated stimulus announced by the central bank but not yet reneged on.

The experiment highlights that incentives to renege on the policy are nontrivial, especially if the shock proves to be of relatively short duration. Of course the benefits of reneging must be weighed against the reputational loss. These considerations are particularly salient in a low-inflation and low-natural interest rate environment where central banks are likely to require forward guidance communication.

6.3 The costs of delay with unanchored expectations

Different country experiences with the zero lower bound on nominal interest rates raise important questions about forward guidance policy. Japan suffered a recession in the early 1990s at a time before many advanced countries made significant changes to monetary policy frameworks by adopting inflation targeting. Consistent with this, Carvalho, Eusepi, Moench, and Preston (2019) show that long-term inflation expectations were poorly anchored, displaying significant instability in response to the down turn. Moreover, the Bank of Japan did not implement unconventional monetary policy until the early 2000s in the form of quantitative easing, and then later still forward guidance. In contrast, when the United States faced the challenges of a zero interest rate environment in the aftermath of the financial crisis or the recent pandemic, they had a well-established highly credible monetary policy regime in which long-term inflation expectations were well anchored. They were quick to introduce zero interest rate policies and were prepared to experiment with the way in which they communicated their forward guidance policy.

We now use the model to explore the consequences of delay in the implementation of forward guidance policy as well as the additional complications that arise from expectations being poorly anchored. Figure 9 reports the results from a central bank that implements optimal forward guidance policy but makes these commitments at progressively later dates. The black lines correspond to the optimal policy implemented at the time of the shock. The red and blue lines the outcomes from optimal forward guidance policy implemented in quarter 10 and 19 respectively. The left column has outcomes for beliefs with benchmark persistence; and the right column unit root beliefs, a proxy for poorly anchored expectations. In each column the successive panels report data on output gap expectations, inflation expectations and the policy promises. For expectations we plot the contemporaneous beliefs at each point in time along with the beliefs at that time about economic conditions five years ahead.

Start with the policy promises for benchmark persistence in beliefs in the left column. Because the figure reports the promised zero interest rate policy conditional on the indicated shock duration, to compute the total time the economy is at the zero lower bound we must add the promise to the shock duration. As an example, ten quarters after the shock hits the economy, the policy promise under the policy enacted in period 1 commits to 7 quarters of zero interest policy (the black bar). In contrast the policy first enacted in quarter 10

Figure 9: Optimal Policy: Is it worth the wait?

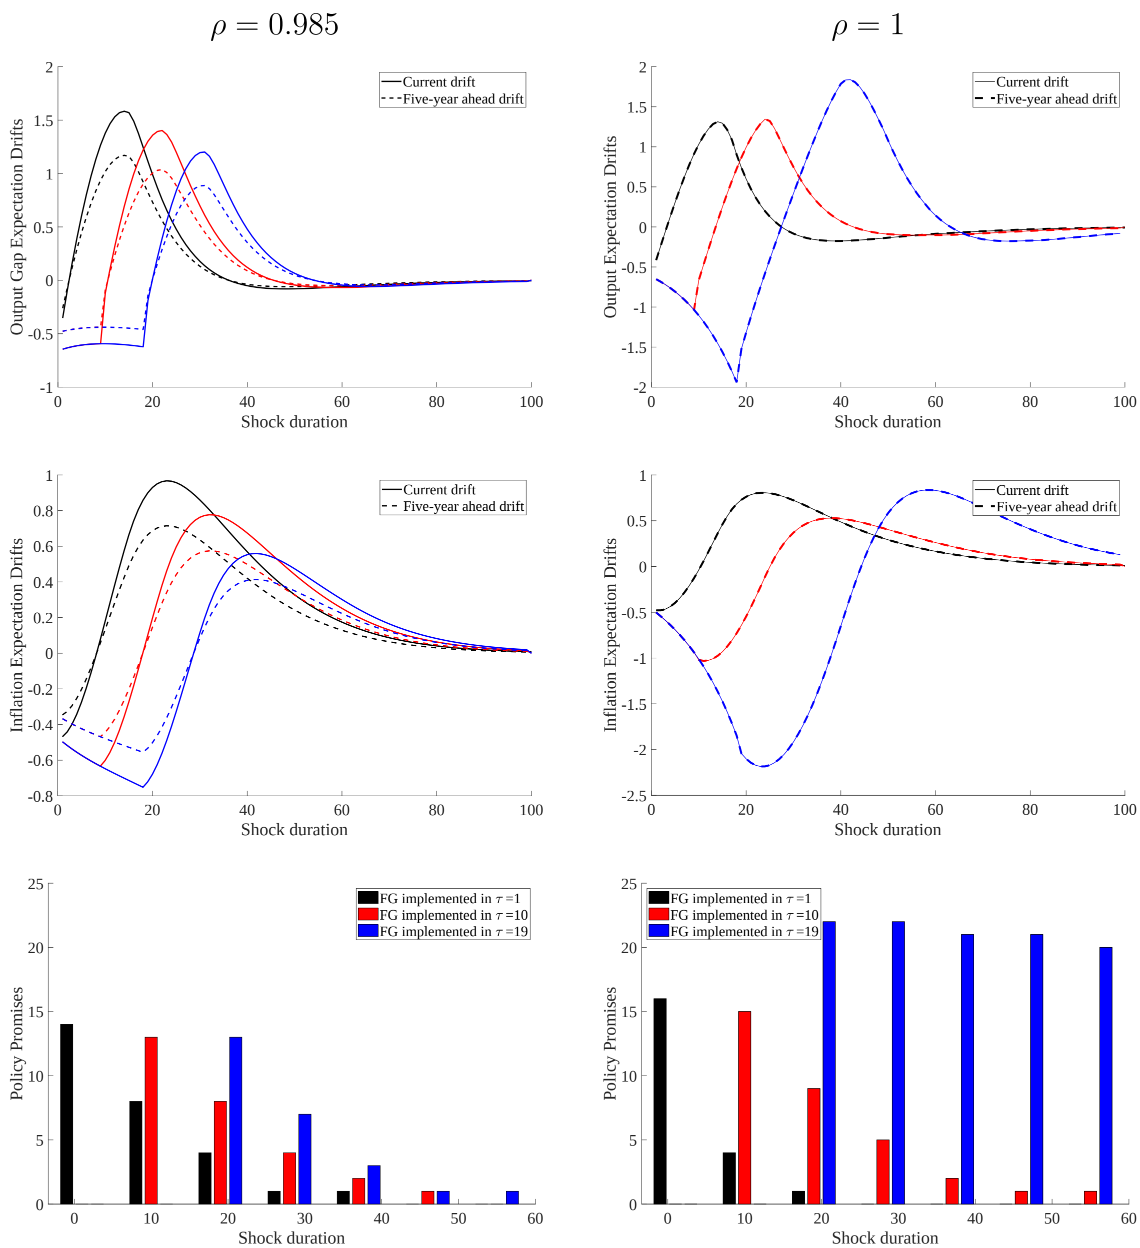

Notes: Shared parameter values $\lambda_x = 0.06$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.024$, $r_L = -0.005$, and $g = 0.075$.

promises 13 quarters of zero interest rates (the red bar). Hence, conditional on a shock having at least duration of 10 quarters, the economy will have 17 and 23 quarters at the zero lower bound respectively. Similar calculations for longer duration shocks, particularly in the case that forward guidance policy is only implemented in quarter 19, reveal that delay has significant consequences for the expected time at the zero lower bound.

The panels above the policy promises plot the current drift and the five-year-ahead expectations held at that time for the output gap and inflation. They make clear that delay leads to falling expectations which makes the constraint of the zero lower bound more acute. Inflation expectations progressively fall until arrested by the implementation of forward guidance policy. Importantly, the five-year-ahead expectations continue to be depressed even after the implementation of forward guidance policy. It is this sluggish response of long-term inflation expectations that explains the increasing time at the zero lower bound when forward guidance is delayed.

The output gap expectations are similar, though initially rise with delay, before falling in the case of a policy implementation delay of 19 quarters. The initial rise in expectations is explained by the stimulatory effects from the rational component of beliefs. With delay the learning component eventually comes to dominate. This also explains the progressive fall in the peak of the hump-shaped dynamics in each profile. The appendix provides analysis of a model which eliminates the stimulatory effect from the rational component of beliefs by assuming that the normalizing constant in beliefs is a fixed constant. The results are similar.

Figure 10: Optimal Policy: Is it worth the wait?

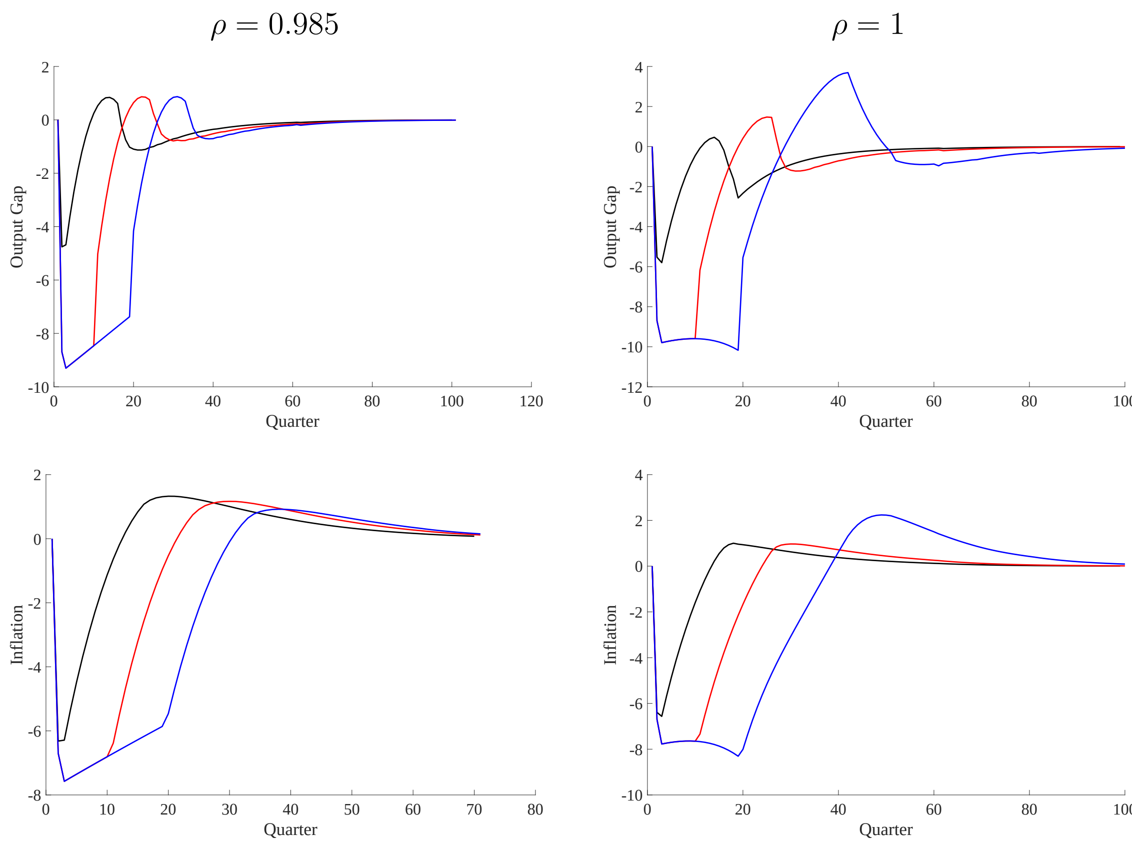

Notes: Shared parameter values $\lambda_x = 0.06$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.024$, $r_L = -0.005$, and $g = 0.075$.

In the case of highly persistent beliefs in the right panel, the effects of drifting expectations are further pronounced (note the different scale and that current and five-year-ahead beliefs are identical when $\rho = 1$). Both output gap and inflation beliefs fall substantially until the implementation of forward guidance policy. And in the case of implementation in quarter 19, inflation expectations continue to fall even after the announcement of forward guidance policy. As a consequence, the required forward guidance ratchets up considerably.

Figure 10 shows the associated outcomes for the output gap and inflation. Because of the implementation delay, the economy experiences a prolonged and large contraction. The contraction is more severe for persistent beliefs. In the case of mean-reverting beliefs in the left column, the optimal policy delivers progressively less over-shooting in inflation as the implementation delay lengthens. Correspondingly, there is not much of a rise in the output boom, and a smaller subsequent recession. In the case of persistent beliefs in the right panel, a larger output boom is required to drag up inflation expectations. As a result, a downturn is still required to later restrain inflation and inflation expectations.

7 Conclusions

This paper determines the optimal forward guidance policy in a model in which long-term expectations can become unanchored and drift downwards. The optimal policy features large front-loaded promises, displaying an insurance principle: aggressive forward guidance stabilizes expectations in the case of persistent shocks but is too stimulatory in the case of transitory shocks. This state-contingent policy is well-approximated by a calendar-based forward guidance policy. A corollary of the insurance principle is the 'too little, too late' principle: because falling long-term expectations mitigate the effects of forward guidance policy, any delay compromises macroeconomic stabilization. In general, macroeconomic demand management is more difficult. By organizing thinking about unanchored inflation expectations; the framing of forward guidance policy announcements; and the merits of aggressive, immediate stimulus and the potential cost of over-heating the economy, the paper informs contemporary policy debate.

References

Andrade, P., and F. Ferroni (2021): "Delphic and odyssean monetary policy shocks: Evidence from the euro area," *Journal of Monetary Economics*, 117(C), 816–832.

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon (2019): "Forward Guidance and Heterogeneous Beliefs," *American Economic Journal: Macroeconomics*, 11(3), 1–29.

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2020): "Imperfect Macroeconomic Expectations: Evidence and Theory," NBER Working Papers 27308, National Bureau of Economic Research, Inc.

Angeletos, G.-M., and C. Lian (2018): "Forward Guidance without Common Knowledge," *American Economic Review*, 108(9), 2477–2512.

Arad, A., and A. Rubinstein (2012): "The 11-20 Money Request Game: A Level-k Reasoning Study," *American Economic Review*, 102(7), 3561–73.

Beechey, M. J., B. K. Johannsen, and A. T. Levin (2011): "Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area Than in the United States?," *American Economic Journal: Macroeconomics*, 3(2), 104–129.

Bern, R., F. Caselli, F. Giglioli, B. Gruss, and W. Lian (2018): "Expectations' Anchoring and Inflation Persistence," Imf working papers, IMF.

Bilbiie, F. O. (2018): "Monetary Policy and Heterogeneity: An Analytical Framework," CEPR Discussion Papers 12601, C.E.P.R. Discussion Papers.

Bodenstein, M., J. Hebden, and F. Winkler (2019): "Learning and Misperception: Implications for Price-Level Targeting," Finance and economics discussion papers, Board of Governors of the Federal Reserve System.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018): "Over-reaction in Macroeconomic Expectations," NBER Working Papers 24932, National Bureau of Economic Research, Inc.

Bosch-Domènech, A., J. G. Montalvo, R. Nagel, and A. Satorra (2002): "One, Two, (Three), Infinity, ...: Newspaper and Lab Beauty-Contest Experiments," *American Economic Review*, 92(5), 1687–1701.

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004): "A Cognitive Hierarchy Model of Games*," *The Quarterly Journal of Economics*, 119(3), 861–898.

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012): "Macroeconomic Effects of Federal Reserve Forward Guidance," *Brookings Papers on Economic Activity*, 43(1 (Spring), 1–80.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2019): "Anchored Inflation Expectations," unpublished manuscript, University of Melbourne.

Cochrane, J. H. (2017): "The New-Keynesian Liquidity Trap," *Journal of Monetary Economics*, 92, 47–63.

Crump, R. K., S. Eusepi, and E. Moench (2011): "Making a Statement: How Did Professional Forecasters REact to the August 2011 FOMC Statement?," Liberty Street Economics, Federal Reserve Bank of New York Blog.

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

Duffy, J., and R. Nagel (1997): "On the Robustness of Behaviour in Experimental 'Beauty Contest' Games," *The Economic Journal*, 107(445), 1684–1700.

Eggertsson, G., and M. Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, pp. 139–211.

Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins (2019): "A Model of Secular Stagnation: Theory and Quantitative Evaluation," *American Economic Journal: Macroeconomics*, 11(1), 1–48.

Eusepi, S. (2010): "Central Bank Communication and the Liquidity Trap," *Journal of Money, Credit and Banking*, 42(2-3), 373–397.

Eusepi, S., M. Giannoni, and B. Preston (2019): "On The Limits of Monetary Policy," unpublished, University of Texas, Austin.

Eusepi, S., and B. Preston (2010): "Central Bank Communication and Macroeconomic Stabilization," *American Economic Journal: Macroeconomics*, 2, 235–271.

——— (2011): "Learning the fiscal theory of the price level: Some consequences of debt-management policy," *Journal of the Japanese and International Economies*, 25(4), 358–379.

Evans, G. W., C. Gibbs, and B. McGough (2019): "A Unified Model of Learning to Forecast," Working paper, University of Sydney.

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

Evans, G. W., and S. Honkapohja (2010): "Expectations, Deflation Traps and Macroeconomic Policy," in *Twenty Years of Inflation Targeting: Lessons Learned and Future Prospects*, ed. by D. Cobham, O. Eitrheim, S. Gerlach, and J. F. Qvigstad, pp. 232–260.

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

Fuster, A., D. Laibson, and B. Mendel (2010): "Natural Expectations and Macroeconomic Fluctuations," *Journal of Economic Perspectives*, 24(4), 67–84.

Gabaix, X. (2020): "A Behavioral New Keynesian Model," *American Economic Review*, 110(8), 2271–2327.

Gali, J. (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction of the New Keynesian Framework*. Princeton University Press.

García-Schmidt, M., and M. Woodford (2019): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," *American Economic Review*, 109(1), 86–120.

Gibbs, C. G., and M. Kulish (2017): "Disinflations in a model of imperfectly anchored expectations," *European Economic Review*, 100, 157–174.

Gurkaynak, R. S., A. Levin, and E. Swanson (2010): "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., UK, and Sweden," *Journal of the European Economic Association*, 8(6), 1208–1242.

Ho, T.-H., C. Camerer, and K. Weigelt (1998): "Iterated Dominance and Iterated Best Response in Experimental "p-Beauty Contests"," *The American Economic Review*, 88(4), 947–969.

Kozicki, S., and P. A. Tinsley (2001): "Shifting endpoints in the term structure of interest rates," *Journal of Monetary Economics*, 47(3), 613–652.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The power of forward guidance revisited," *American Economic Review*, 106(10), 3133–58.

McKay, A., E. Nakamura, and J. Steinsson (2017): "The Discounted Euler Equation: A Note," *Economica*, 84(336), 820–831.

NAGEL, R. (1995): "Unraveling in Guessing Games: An Experimental Study," *The American Economic Review*, 85(5), 1313–1326.

ORPHANIDES, A., AND J. C. WILLIAMS (2008): "Learning, expectations formation, and the pitfalls of optimal control monetary policy," *Journal of Monetary Economics*, 55(Supplemen), S80–S96.

PRESTON, B. (2006): "Adaptive Learning, Forecast-Based Instrument Rules and Monetary Policy," *Journal of Monetary Economics*, 53, 507–535.

SARGENT, T. J., AND N. WILLIAMS (2005): "Impacts of Priors on Convergence and Escape Dynamics," *Review of Economic Dynamics*, 8(2), 360–391.

WIEDERHOLT, M. (2015): "Empirical Properties of Inflation Expectations and the Zero Lower Bound," unpublished.

WILLIAMS, J. C. (2010): "Monetary policy in a low inflation economy with learning," *Economic Review*, pp. 1–12.

WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

——— (2018): "Monetary Policy Analysis when Planning Horizons are Finite," NBER Working Papers 24692, National Bureau of Economic Research, Inc.

WOODFORD, M., AND Y. XIE (2019): "Policy Options at the Zero Lower Bound When Foresight is Limited," *AEA Papers and Proceedings*, 109, 433–37.

A details and Derivations

[TO BE ADDED]