Monetary Policy Trade-offs at the Zero Lower Bound*

Stefano Eusepi† Christopher G. Gibbs‡ Bruce Preston§

Abstract

We study zero interest-rate policy in response to a large negative demand shock when long-run inflation 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. Optimal policy is well-approximated by a constant calendar-based forward guidance, independent of the shock's realized persistence. This insurance principle qualitatively and quantitatively distinguishes our paper from other recent research on bounded rationality and the forward guidance puzzle.

Keywords: Optimal Monetary Policy, Learning Dynamics, Expectations Stabilization, Forward Guidance
JEL Codes: E32, D83, D84

*July 2, 2024. The authors thank seminar and conference participants at the Bank of Finland, SNDE Meetings 2019, Reserve Bank of Australia, the Virtual Australian Macroeconomics Seminar, the Virtual East Asia Macroeconomics Seminar, 2021 Conference on "Expectations in Dynamic Macroeconomic Models", University of Washington, Northwestern University and particularly Larry Christiano, Chris Edmond, Greg Kaplan, Eric Leeper and our discussant Nigel McClung. 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.

> The historical record is thick with examples of underdoing it ... And pretty much in every cycle, we just tend to underestimate the damage and underestimate the need for a response. I think we've avoided that this time.
>
> — Jerome Powell (2021)

1 Introduction

From two large macroeconomic events—the global financial crisis and the pandemic—come two opposing views on the merits of stabilization policy. One claims that aggressive monetary stimulus prevented deflation, with all the devastating consequences of the Great Depression of the 1930s. The other claims that same stimulus to be overly inflationary, inviting a return to the misfortunes of the Great Inflation of the 1970s. This paper argues that these views are both essential ingredients of optimal monetary policy. They each reflect aspects of a general trade-off that confronts monetary policy at the zero lower bound on interest rates, which we call the *insurance principle*.

Two properties of macroeconomic reality reconcile these views. First, large shocks and new policies make past experience a poor guide to future developments, so that people must learn from recent experience. Second, central banks must stabilize the economy in response to a shock with uncertain duration. By gradually propagating and amplifying the effects of shocks and policy over time, learning exacerbates the well-known tension between optimal "commitment" and "discretionary" monetary policy. There can be large differences between policies that are optimal ex ante and those that are optimal ex post, after observing the realized persistence of the shock. We contend that the "deflationary" and "inflationary" views of monetary policy in response to large negative demand shocks reflect, respectively, ex ante and ex post perspectives about what constitutes desirable policy.[1]

In a new Keynesian model, we show the optimal forward guidance policy is a set of front-loaded state-contingent promises of future zero interest rates. Commitments are initially large, lasting some several years, but decline over time if the shock persists. This structure of promises means calendar-based forward guidance policies well-approximate the fully optimal policy.[2] In the case of a persistent shock, this structure of promises prevents self-fulfilling deflation. The policy provides short-run support of long-term inflation expectations which maintains the efficacy of monetary policy when the central bank confronts an unfavorable

[1] To be clear, in referencing this debate about the appropriate stance of monetary policy, we are neither arguing the global pandemic is reducible to a demand shock nor that fiscal policy is unimportant.

[2] Calendar-based policies promise a fixed period of zero interest rate policy rather than a set of state-contingent promises. Depending on the specific environment, we find promises of three-to-five years are close to optimal for the size of demand shock we consider.

persistent shock. Of course, in the event of a favorable transitory shock, fulfilling past commitments to a sustained period of zero interest rates will be too stimulatory. The policy generates a substantial over-shooting of the inflation target, requiring an economic downturn to restrain inflation expectations. The central bank optimally accepts poorer stabilization of inflation in response to transitory shocks to eliminate deflationary expectations and achieve better stabilization when shocks are more persistent. This trade-off is the insurance principle.

The results are strikingly different to Eggertsson and Woodford's (2003) full-information rational expectations analysis, where this trade-off is absent. Optimal policy under these assumptions prescribes a profile of state-contingent promises that are back-loaded: commitments to zero interest rates are initially small—just one quarter for short-duration shocks—and rise only modestly with the duration of the shock. This policy almost completely stabilizes the macroeconomy in response to a negative demand shock. We show these differences reflect the fact that general equilibrium effects of forward guidance policy under rational expectations are 'front-loaded', occurring immediately at the time of the announced policy, but under learning are 'back-loaded', occurring with a delay.

Central to our results are boundedly rational households and firms that have imperfect information about the true data generating process describing macroeconomic outcomes. Agents must learn about how the shock and policy response affects the economy. They revise their expectations by extrapolating from recent forecast errors. An unanticipated negative shock causes agents to mark down their inflation and output gap beliefs. Because these mark-downs affect subsequent dynamics and beliefs, the initial shock gets propagated over time and expectations become partially self-fulfilling.

When a central bank cannot credibly commit to forward guidance and must use short-term interest rate policy the economy experiences a deflationary spiral while the demand shock persists. Learning makes the decline in real activity larger than in a rational expectations economy, with deeper recessions associated with more persistent realizations of the demand shock. The economy spends a substantially longer time at the zero lower bound. Conventional monetary policy is unable to arrest deflationary pressure from falling inflation expectations which raise real interest rates and make the zero lower bound a more significant constraint.

We next consider a central bank that can make credible commitments to future zero interest rate policy. Agents perfectly understand the implications of forward guidance for the path of future interest rates. Consistent with our assumption that agents must learn about the macroeconomic consequences of the shock and policy, they cannot evaluate the general equilibrium implications of forward guidance policy. Expectations about future output and inflation are only revised subsequently, in response to changing economic conditions. This property is consistent with empirical evidence on the effects of forward guidance on professional forecaster expectations adduced by Crump, Eusepi, and Moench (2011), Del Negro, Giannoni, and Patterson (2012), Campbell, Evans, Fisher, and Justiniano (2012) and Andrade and Ferroni (2021).[3]

To overcome weaker general equilibrium effects, the optimal forward guidance policy makes large front-loaded commitments to zero interest rate policy. This prevents a deflationary spiral and supports long-term inflation expectations in the event of highly persistent shocks. This policy provides nominal space in the event of adverse shocks by widening the gap between the expected interest rate and its natural level, and, therefore, scope for further monetary support. But in the event of transitory shocks requires a substantial over-shooting of the inflation target—offering a possible rationalization of the currently high inflation in the United States as the outcome of optimal policy under an unexpectedly favorable shock. In this way optimal policy provides self insurance (at the level of the economy) against the risk of a demand shock with unknown duration.[4]

A special case of our model in which prices are fixed provides further understanding and intuition for the insurance principle. We decompose analytically the output response to forward guidance into two effects: the partial equilibrium effect of policy (the change in demand given a change in the future interest rate path, holding fixed current and expected future output) and the general equilibrium effect (the adjustment in output from market clearing and subsequent dynamics from learning).

We show that the partial equilibrium effects of policy under learning are identical to a rational expectations analysis. In contrast, the general equilibrium effects are different. On impact, at the date of the forward guidance announcement, the general equilibrium effects of policy are dramatically smaller than a rational expectations analysis. This reflects the fact that output expectations are initially insensitive to the announced interest rate path. In this way, our model is similar to other resolutions of the forward guidance puzzle, such as Angeletos and Lian (2018), Farhi and Werning (2019), Gabaix (2020) and Bilbiie (2018), which all make assumptions to moderate the strength of general equilibrium responses to policy announcements. What distinguishes our work from these papers is that learning propagates and amplifies these general equilibrium effects over time. That is, the general equilibrium effect is dynamic, rather than static as in these and rational expectations models.

[3]This approach to implementing central bank communication was earlier introduced by Preston (2006) and Eusepi and Preston (2010) and is an example of level k reasoning in which agents perform a single round of deductive reasoning in response to the policy announcement. And is therefore related to Farhi and Werning (2019) which explores different levels of deductive reasoning for the forward guidance puzzle.

[4]To complete the analogy to insurance, the premium, worse stabilization outcomes for transitory shocks, secures payments in the form of better stabilization outcomes for persistent shocks.

The dynamic general effect on aggregate demand can grow to be substantial as long-term expectations adjust over time, much larger than in a rational expectations analysis.

This decomposition sheds light on the structure of optimal forward guidance promises. Initially weak general equilibrium effects reduce the effective stimulus from forward guidance policy, requiring larger promises than under rational expectations. However, these aggressive zero interest rate commitments increase inflation expectations and lower real interest rates, further raising inflation expectations. Self-fulfilling movements in long-term expectations lead to macroeconomic effects which build gradually over time and are persistent—beyond the period of the negative demand shock and zero interest rate policy. The dynamic general equilibrium effects, which generate a boom in real activity, explain why there is a trade-off and why promises are tapered. Because general equilibrium effects are 'back-loaded' forward guidance must be 'front-loaded'. In contrast, under rational expectations, the general equilibrium effects are 'front-loaded' making forward guidance 'back-loaded'.

We provide three applications of the insurance principle. The first compares our findings to Gabaix (2020) and Bilbiie (2018) who also propose models that resolve the forward guidance puzzle. We show that the predictions of these models are qualitatively similar to the optimal commitment rational expectations model. They imply the optimal forward guidance policy is back-loaded, with promises rising with the duration of the shock.

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 later renege on these promises. The value of reneging is largest in the case of favorable short-duration shocks, because raising interest rates early limits the extent to which higher inflation becomes entrenched in inflation expectations, requiring much less restraint in real activity. The experiment highlights the fundamental trade-off manifest in the ex ante and ex post perspectives on aggressive stabilization policy.[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 advanced 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 the 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

[5] Of course, this is not 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.

guidance 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 implementation of policy is delayed, the trade-offs become increasingly less favorable requiring even more aggressive forward guidance policy.

Related literature. The stability of long-term expectations are central to our results. Using a structural model, Carvalho, Eusepi, Moench, and Preston (2019) provide evidence, for the United States, Japan and other countries, 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. Bern, Caselli, Giglioli, Gruss, and Lian (2018) provides a comprehensive measure of the degree of anchoring for a large set of countries.

Bounded rationality and imperfect information have implications for monetary policy design. Ferrero (2007), Gaspar, Smets, and Vestin (2006), Orphanides and Williams (2005) and Molnar and Santoro (2013), show optimal monetary policy prescribes a more aggressive response to inflation and muted activism toward output gap stabilization compared to a rational expectations model. Closer to the current paper, Eusepi, Giannoni, and Preston (2019) show that active demand management is not desirable when monetary policy has imprecise control of long-term interest rate expectations. Gáti (2021) confirms these insights in a model with an endogenous gain. Gibbs and Kulish (2017) show that unanchored expectations complicates the design of disinflation policy and raises costs. But none of these papers study the zero lower bound or, more broadly, the use forward guidance as an additional policy instrument.

In a new Keynesian model with learning Williams (2010) shows Taylor-type rules responding to the price level are undesirable when dealing with the zero lower bound, despite performing well in models with rational expectations. Using global solution methods Eusepi (2010), Evans, Guse, and Honkapohja (2008) and Evans and Honkapohja (2010) find that unstable expectations can result a prolonged period of low or negative inflation and low levels of output, close to a liquidity trap equilibrium. These authors show central bank transparency about policy rules improve stabilization outcomes, but do not analyze forward guidance. Similar to our paper, Evans, Honkapohja, and Mitra (2021) studies the effects of large fiscal and monetary expansions on stabilizing expectations in this nonlinear model environment. However, optimal state-contingent 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), Gabaix (2020), and Gibbs and McClung (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) show that imperfect common knowledge of forward guidance announcements can also mute general equilibrium effects and therefore limit its impact on the economy. Similar to our paper, they show that front-loading of fiscal policy announcement is desirable, although the mechanism is different. However, we show that optimal monetary policy in many of these models is similar to the full information benchmark. Importantly, these papers do not address learning about monetary policy, 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 to weaker general equilibrium effects induced by bounded rationality and learning.

2 The Model

We use a simple New Keynesian based on Woodford (2003) and Gali (2008). This facilitates analytical results and comparison to other recent papers on this topic. Aggregating a log-linear approximation to the optimal individual consumption and pricing decisions of households and firms gives the demand and supply equations

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

$$\pi_t = \hat{E}_t \sum_{T=t}^{\infty} \left(\xi\beta\right)^{T-t} \left[\kappa x_T + (1-\xi)\, \beta \pi_{T+1}\right] \tag{2}$$

where x_t denotes the output gap; r_t^n the natural rate of interest, an exogenous process; π_t the inflation rate; and R_t the short-term nominal interest rate.[6] The operator $\hat{E}_t$ denotes average subjective expectations, which might differ from rational expectations.

The aggregate demand equation determines the output gap as the discounted expected value of future income, with discount rate $0 < \beta < 1$, where the second term captures the effect of variations in the real interest rate, applied in future periods, due to changes in the nominal interest rate and goods price inflation (with an intertemporal elasticity of substitution of $\sigma > 0$). The aggregate supply curve determines inflation as the discounted future sequence of output gaps and the inflation rate, which also depends on the exogenous

[6] The appendix offers a detailed derivation of the model.

probability $0 < \xi < 1$ of not being able to reset their price in any subsequent period. The slope of the Phillips curve is measured by $\kappa = (1 - \xi\beta)(1 - \xi)/\xi$.

2.1 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.

Information. Agents are boundedly rational and have imperfect information about the true data-generating processes of aggregate variables they need to forecast to make spending and pricing decisions. They do not know the structural equations (1) and (2), which are implications of the aggregation of individual optimal decision rules and market clearing. Before describing how 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 zero lower bound 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, as a result of the large shock, uncertain about the equilibrium level of output and inflation in both the low and high 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. Expectations about the policy rate depend on the expected fundamentals of the economy—the natural rate of interest and inflation expectations—and on the policy announcements.

Subjective expectations. The economy starts in the low state where $r_t^n = r_L < 0 < r_H$. In each successive period the natural rate reverts back to the high state r_H with a constant probability $0 < \delta < 1$. Given this information, agents form estimates of equilibrium inflation and output gap based on the forecasting model

$$z_t = z_S + \bar{\omega}_t + e_t \tag{3}$$

$$z_S = \Omega_S r_S \tag{4}$$

$$\bar{\omega}_{t+1} = \rho\bar{\omega}_t + u_{t+1} \tag{5}$$

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

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

where $0 \leq \rho \leq 1$ and $\Omega_S > 0$ are parameters; e_t and u_t i.i.d. with $J = \hat{c}Q$ for $J = E\left[e_t e_t'\right]$, $Q = E\left[u_t u_t'\right]$ and scalar $\hat{c}$.

The forecasting model has two components. The first is a state-contingent vector of constants, z_S, with dependence on the observed natural interest rate determined by the parameters Ω_S. These parameters capture time-invariant prior beliefs about the consequences of a negative demand shock. The second is the sum of two unobserved factors, a persistent drift, $\bar{\omega}_t$, and an i.i.d. noise disturbance, e_t. These variables measure residual uncertainty about the data generating process z_t, including economic effects from changes in the natural rate that are not fully captured by z_S or the effects of policy. The parameter ρ captures the persistence of the drift and the scalar $\hat{c}$ measures the volatility of the drift innovation, u_t, relative to that of e_t. The latter captures how informative the data are about $\bar{\omega}_t$. Priors about these parameters are time-invariant.

Estimates of the unobserved drift given information available up to time t, $\omega_{t+1|t}$, are updated each period using the Kalman filter

$$\omega_{t+1|t} = \rho\omega_{t|t-1} + \rho\gamma\left(z_t - z_S - \omega_{t|t-1}\right) \tag{6}$$

where $0 < \gamma < 1$ is a function of the parameters ρ and $\hat{c}$.[7] The evolution of $\omega_{t+1|t}$ captures how agents gradually learn about the effects of the shocks or policy announcements.[8] For example, a positive inflation surprise leads to an upward revision to the estimate of the inflation drift. The learning gain $g \equiv \rho\gamma$ measures the sensitivity of agents' estimates to recent forecast errors and it is directly related to the drift's persistence and the signal to noise ratio.

The forecasting model (3)-(5) together with the estimated drift imply conditional expectations for output gap and inflation in the low state satisfy

$$\hat{E}_t(z_{T+1}|S = L) = (1-\delta)^{T+1-t} z_L + \left(1 - (1-\delta)^{T+1-t}\right) z_H + \rho^{T-t}\omega_{t|t-1}. \tag{7}$$

The first component determines the expected path as a function of the expected path of the

[7] See the online appendix for a full derivation.

[8] 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 $t-1$. This avoids a complex simultaneity arising from having current beliefs depend on current macroeconomic outcomes.

natural rate of interest, while the second component captures the impact of the estimated drifts, an effect independent of the state. For sufficiently high values of ρ, changes in the estimated drifts have impact on forecast horizons beyond the period in which the economy is in the low state.

We first consider an economy where the central bank cannot use credible announcements to influence expectations. This benchmark introduces the challenges confronting optimal policy and identifies key differences relative to a full-information rational expectations economy. Without forward guidance, interest rate expectations satisfy

$$\hat{E}_t\left(R_{T+1}|S=L\right) = (1-\delta)^{T+1-t}\times 0+\left(1-(1-\delta)^{T+1-t}\right)\left(r_H+\rho^{T-t}\omega^{\pi}_{t|t-1}\right). \quad (8)$$

Conditional expectations of future interest rates reflect two assumptions. First, agents understand interest rates will be zero in the low state. Second, they form expectations of the interest rate in the high state anticipating policy tracks the 'nominal' natural rate of interest

$$\omega^R_{t|t-1}=\omega^{\pi}_{t|t-1}+r_H,$$

which we refer to as the Fisher equation. This implies inflation expectations play an important role in shaping interest rate expectations. While this subjective expectation differs from the true data generating process of the policy rate under the optimal policy, it is consistent with the belief that the central bank optimally lets inflation persistently deviate from a constant target. In steady state, subjective and objective expectations coincide.

A Special Case. Given the two-state Markov process for the natural rate, the aggregate demand and supply equations have a rational expectations solution with $x_H=\pi_H=0$ and $R_H=r_H$ in the high state and

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

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

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.[9]

Throughout the paper we assume the state-contingent constants in beliefs, z_S, are consistent with the rational expectations optimal discretion equilibrium. The above solutions

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

implicitly define the constants Ω_S. 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 smaller than under full information. Conversely, if initial beliefs were pessimistic the equilibrium outcomes would be larger. We center beliefs at rational expectations. The impact response in the two economies is 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.[10]

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 is a measure of how well anchored expectations are.

2.2 Objective Beliefs and Self-fulfilling Deflations

Using the above assumptions to evaluate expectations in (1) and (2), provides the true data-generating process

$$x_t^H = -\sigma(R_t - r_H) + \frac{1-\beta}{1-\beta\rho}\omega^x_{t|t-1} + \frac{\sigma(1-\beta)}{1-\beta\rho}\omega^\pi_{t|t-1} \tag{11}$$

$$\pi_t^H - \kappa x_t^H = \frac{(1-\xi)\beta}{1-\xi\beta\rho}\omega^\pi_{t|t-1} + \frac{\kappa\xi\beta}{1-\xi\beta\rho}\omega^x_{t|t-1} \tag{12}$$

in the high state and

$$x_t^L = x_L + \frac{1-\beta}{1-\beta\rho}\omega^x_{t|t-1} + \sigma\left[\frac{1-\beta}{1-\beta\rho} + \frac{\beta\,(1-\delta)}{1-\beta\rho\,(1-\delta)}\right]\omega^\pi_{t-1} \tag{13}$$

$$\pi_t^L - \kappa x_t^L = \pi_L - \kappa x_L + \frac{(1-\xi)\beta}{1-\xi\beta\rho}\omega^\pi_{t|t-1} + \frac{\kappa\xi\beta}{1-\xi\beta\rho}\omega^x_{t|t-1} \tag{14}$$

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.

[10] Section 5 and the online appendix show this assumption does not affect our qualitative policy conclusions. It only serves to scale the initial shock faced by the policymaker.

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 $\omega^x_{t|t-1}$), the substitution effects from interest rate policy differ (the terms in $\omega^\pi_{t|t-1}$). 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), the low state equilibrium output gap under rational expectations. The elasticity with respect to expected future real rates is also 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.

Substituting the expressions for equilibrium output and inflation, (13) and (14), into the belief updating equations, (6), gives the true data-generating process. Belief dynamics are a complicated function of the parameters ρ and γ, and the probability of exiting the bad state, δ. To ease exposition, let us consider the case $\rho = 1$, which delivers the following system

$$\begin{aligned}
\omega^x_{t+1|t} &= \omega^x_{t|t-1} + \gamma\sigma\left[\frac{1}{1-\beta(1-\delta)}\right]\omega^\pi_{t|t-1} \\
\omega^\pi_{t+1|t} &= \gamma\left(\frac{\kappa}{1-\alpha\beta}\right)\omega^x_{t|t-1} + \left(\kappa\sigma\left[\frac{1}{1-\beta(1-\delta)}\right] + \frac{(1-\alpha)\beta}{1-\alpha\beta}\right)\omega^\pi_{t|t-1}.
\end{aligned}$$

Falling output expectations mark down future output beliefs one for one. But they also lead to a mark down in inflation expectations which feeds back into a further decline in output expectations. Moreover, a higher learning gain or a longer expected duration of the zero lower bound period has further de-stabilizing effects. Without forward guidance, this dynamic can be unstable. The appendix proves the following result.

Proposition 1. *Consider the system defined by equations (13)-(14) and (6). There exists a persistence parameter $0 < \rho^* \leq 1$ such that for $\rho > \rho^*$ the system has at least one eigenvalue strictly outside the unit circle for all maintained parameter values.*

A central model property then is that for sufficiently persistent subjective beliefs, the equilibrium dynamics of beliefs at the zero lower bound will be locally explosive. Through falling future expected income and rising future expected interest rates, pessimism induces strong negative general equilibrium effects which weigh down on aggregate demand. We later show a key implication from the proposition (and a crucial difference with rational expectations) is the link between shock duration and inflation beliefs. As the duration of the shock lengthens, the deterioration in long-term inflation expectations is more marked the more persistent are beliefs. Moreover, higher gains represent expectations that are progressively less well anchored. This properties are critical to our results on forward guidance.

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

$$L_t = E_t \sum_{T=t}^{\infty} \beta^{T-t} \left(\pi_T^2 + \lambda_x x_T^2 \right) \tag{15}$$

where $0 < \beta < 1$ and $\lambda_x > 0$ 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.

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 -1.2 percent per annum. The natural rate reverts back to the stead-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]

What about initial beliefs? At the time of the shock (period 1 in our experiment) we assume the economy is at the rational expectations equilibrium and so $\omega^x_{1|0} = \omega^\pi_{1|0} = 0$. We then set agents' initial estimate in period 2 (after the shock) as a function of the output and inflation correction on impact

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

With this choice, we make the impact effect of the shock identical under learning and rational expectations. Subsequent equilibrium outcomes differ only because of the endogenous propagation of this forecast error. Beliefs in periods $t > 2$ then satisfy (6).

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

Figure 1: Optimal Policy without Forward Guidance

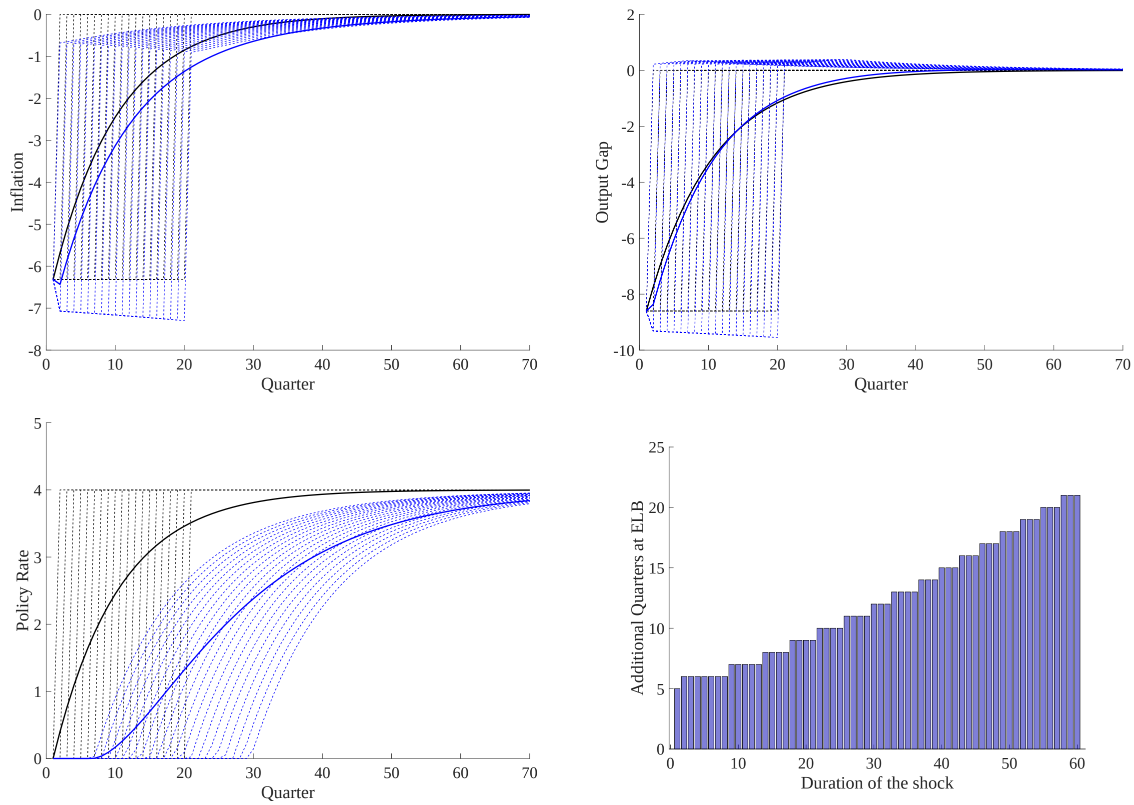

Notes: RE discretion (black) compared to discretion with learning (blue). Parameter values $g = 0.075$, $\rho = 0.985$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\lambda_x = 0.05$, $\kappa = 0.02$, $r_L = -0.003$.

3.3 Benchmark Results

Three panels in Figure 1 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 clarity, 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 the true data-generating process. These paths therefore summarize the dynamics that the central bank expects at the time of the shock.[12] 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

[12] Formally, these are calculated using analogues to equation (7).

returns to steady state: monetary policy raises interest rates when it is feasible, and output and inflation are completely stabilized. The sequence of realizations for output and inflation along these black trajectories are the constants appearing in agent's beliefs under learning.

Comparing the dynamics under optimal policy with learning grants four insights. First, consistent with Eusepi and Preston (2011), learning 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, the more persistent the shock, the deeper the recession. 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. Fourth, 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.[13]

3.4 The Effects of Expectations Anchoring

Proposition 1 shapes these quantitative results. Conditional on being in the low state, self-fulfilling deflations are possible because monetary policy is constrained. More poorly anchored expectations deliver prolonged periods at the zero lower bound. We interpret long-term expectations that are both more sensitive to short-run forecast errors and more persistent as being less well anchored.

Figure 2 plots outcomes for different belief parameters in response to the natural rate shock. The top row gives the five-to-ten year average inflation expectation conditional on reaching each duration of the shock. The bottom row gives the additional time at the zero lower bound. 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.

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. As the duration of the shock lengthens, the deterioration in long-term inflation expectations is more

[13] 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.

Figure 2: The Stability of Long-term Expectations

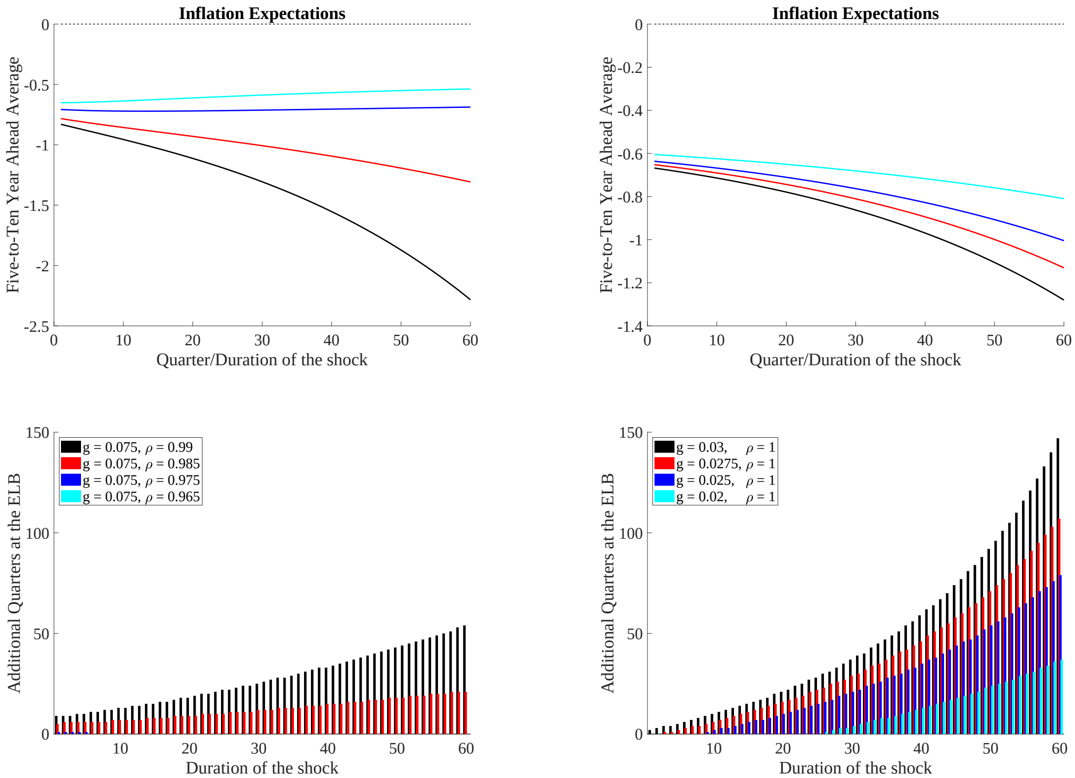

Notes: Top row of figures shows the evolution of the average of five- and ten-year ahead inflation expectation held by agents in quarter t conditioning on the shock lasting until at least quarter t. Parameter values $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\lambda_x = 0.05$, $\kappa = 0.02$, $r_L = -0.003$.

marked the more persistent are beliefs. For beliefs with persistence greater than $\rho^* \approx 0.98$ there is a pronounced downward drift in long-term expectations. Conditional on being in the low state, the equilibrium dynamics of inflation and output beliefs are explosive.[14] Importantly, Crump, Eusepi, and Moench (2016), Carvalho, Eusepi, Moench, and Preston (2019), Eusepi, Giannoni, and Preston (2019) and Crump, Eusepi, Moench, and Preston (2021) adduce empirical evidence that ρ takes values of unity or very near unity. Conversely, beliefs persistence below the threshold ρ^* leads to modestly rising long-term expectations and a dramatic decline in the time at the zero lower bound.[15]

Turning to the right column, with unit root subjective beliefs the equilibrium dynamics of inflation and output beliefs are unstable. 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. Impor-

[14] A value of $\rho = 0.965$ implies a half-life in the drift process just about five years, while a value of $\rho = 0.99$ results in a half-life of nearly twenty years.

[15] For the two smallest values of the persistence parameter the effects are barely visible on the graph.

tantly, this additional time at the zero lower bound is not a deflation trap as is commonly found in papers with adaptive learning.[16] Dynamics in these periods are similar to those shown in Figure 1 with output and inflation persistently, but not significantly, away from steady state.

4 Forward Guidance Policy

Now suppose the central bank can commit to zero interest rate policy in the future. This additional instrument allows policy stimulus while the economy is a the zero lower bound. 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 $\tau \in [1, 2, 3, ...]$. 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.

Forward guidance announcements lead households to revise only their beliefs about future interest rate policy. Because they do not know the structural equations defining the economy, households are unable to evaluate the general equilibrium implications of the new policy for future inflation and output. Consistent with this, the constants, z_S, in the forecasting model (3)-(5), remain fixed at the rational expectations equilibrium in absence of any forward guidance announcement.[17] Expectations about future inflation and the output gap therefore remain unchanged initially. However, these beliefs are revised in future periods in response to changed macroeconomic conditions that result from the announced policy. This is particularly reasonable in situations of the kind we are interested in, such as the economic consequences and policy responses to a global pandemic, about which agents are likely to be uncertain.

This assumption is consistent with evidence from surveys of professional forecaster 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

[16] This conclusion does depend on model parameters. Much larger Kalman gains and longer expected durations of the negative demand shock can lead to a deflation trap.

[17] Our conclusions about the general character of policy do not depend on this assumption. For example, the appendix shows results when the constant is equal to zero.

modest changes to the average term structure of expectations for inflation and output and their distribution across individual forecasters.

4.1 The policy problem

In response to a negative demand shock in period 1, the central bank chooses forward guidance policy to minimize the loss function (15) 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. Following Eggertsson and Woodford (2003) we use a numerical approximation of the optimal policy response to a shock given by a two-state Markov process under which, conditional on the bad state lasting for any length of time, it always has an expected further duration of another 10 quarters.[18]

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 zero lower bound. 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 as described by equation (11). In all regimes, forward guidance policy only affects the structure of the aggregate demand relationship and, in particular, expectations about the nominal rate of interest. 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

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

where $x_t^{k_\tau(j)}$ denotes the output gap in a period under a promise of k_τ periods of zero interest rates with $k_\tau - j$ the remaining periods until lift off. Therefore, $j = 0$ corresponds to the

[18] We assume the natural rate shock reverts to steady state with certainty after 400 periods. Choices of truncation point of 120 periods or larger gives the same results first 60 quarters.

time when the economy switches back to the high state, and the central bank implements k_τ periods of additional zero interest rates.

The first term is standard, the wealth effects from future anticipated income. The second term in brackets captures the intertemporal substitution effects from variations in interest rates. It measures the difference between the sequence of expected future one-period nominal interest rates (the first term in parentheses) and the sequence of expected future one-period 'natural' nominal interest rates, defined as sum of the natural real rate of interest plus inflation expectations (the second term in parentheses). As always in the new Keynesian model, it is this gap which determines the effective stimulus from monetary policy, not the level of nominal of interest rates. For given inflation expectations, the effects of zero interest rate policy in reducing nominal rates relative to the nominal neutral rate is seen in the additional discounting for $k_\tau - j$ periods. Movements in inflation expectations also matter for stimulus. Rising inflation expectations increase the long-term neutral nominal rate granting more nominal space and scope for stimulus. Falling inflation expectations decrease the long-term neutral rate, reducing nominal space and stimulus.[19]

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. The appendix shows aggregate demand can be written as

$$\begin{aligned} x_t^L &= x_L + \frac{(1-\beta)}{1-\beta\rho}\omega_{t|t-1}^x \\ &\quad -\sigma\beta\left(\Psi_t^{\tilde{\rho}}\omega_{t|t-1}^\pi + \Psi_t^1 r_H\right) + \sigma\left(\frac{1}{1-\beta\rho}\omega_{t|t-1}^\pi + \beta\left[\frac{1}{1-\beta} - \frac{1-\delta}{1-\beta(1-\delta)}\right] r_H\right). \end{aligned}$$

where, for $\tilde{\rho} = \{\rho, 1\}$,

$$\Psi_t^{\tilde{\rho}} = (1-\delta)\,\beta\Psi_{t+1}^{\tilde{\rho}} + \delta\frac{(\beta\tilde{\rho})^{k_t}}{1-\beta\tilde{\rho}}$$

is a first-order difference equation which encodes the effects of future commitments to zero interest rates.

To unpack these expressions note that when $\rho = 1$, the expected discounted sequence of

[19] This discussion is normally cast as the gap between the real interest rate relative to the real natural rate of interest. Falling inflation expectations increase real interest rates, lowering the gap between the real and natural rate of interest. Whether nominal or real, the implications are identical.

future interest rates (effectively, a long-term rate) is given by

$$E_t \sum_{T=t}^{\infty} \beta^{T-t} R_{T+1} \quad = \quad \Psi_t^1 (\omega_{t|t-1}^{\pi} + r_H)$$

where $\omega_{t|t-1} + r_H$ is the expected one-period nominal interest rate in each future period, conditional on a return to the high state. Solving the difference equation forward provides

$$\Psi_t^1 \quad = \quad \sum_{j=0}^{\infty} \delta(1-\delta)^j \beta^j \frac{\beta^{k_{t+j}}}{1-\beta}.$$

This expression captures the effect of all future promises on expected interest rates today.

Each element in the sum reflects one particular contingency that takes the economy to lift-off, the resumption of conventional interest rate policy. For example, consider the case when the economy returns to the high state in period j and the central bank implements k_{t+j} periods of zero interest rate policy. This occurs with probability $\delta(1-\delta)^j$. Lift-off then occurs after period $j + k_{t+j}$. At that time, the long-term interest rate is expected to be $(1-\beta)^{-1}(\omega_{t|t-1}^{\pi} + r_H)$ given current beliefs in period t. Of course, to obtain the interest rate faced in period t, in this particular contingency, we simply compute the present value, discounting by $\beta^{j+k_{t+j}}$, since one-period interest rates prior to lift-off are equal to zero. The actual expected long-term interest rate in period t is then the sum of these interest rates in all possible contingencies. State-contingent forward guidance therefore lowers current interest rates by shifting beliefs about what interest rate will apply in all future realizations of uncertainty.

This basic logic continues to apply in the case that $0 < \rho < 1$, but with modified effects from the increased discounting that applies to inflation expectations. In general, the larger the promises, the smaller is $\Psi_t^{\tilde{\rho}}$ and for constant promises k a fixed amount of forward guidance is more effective when implemented in period t rather than $t+1$ because of the effects of discounting. Together this means that earlier and larger action is more stimulatory. The model displays a 'too little, too late' property.

With this understanding, we can interpret the aggregate demand equation. The first term comes from the state-dependent constant in the forecasting model, households' prior understanding of the consequences of a large negative demand shock. The second term wealth effects from anticipated future income. The third term gives the expected sequence of future one-period nominal interest rates given forward guidance. The final term gives the expected sequence of neutral nominal rates, conditional on being in the low state. The expression in brackets, therefore, accounts for the expected duration of the low state. Again, for given inflation expectations, the effective stimulus from forward guidance is given by the difference between the third and final terms.

In contrast to regime 2, the efficacy of zero interest rate policy depends on both inflation expectations and the expected duration of the negative demand shock. As before, falling inflation expectations reduce nominal space and therefore the stimulus from a given forward guidance policy. But so does an increase in the expected duration of the negative demand shock. 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. And adverse shocks, with longer expected durations, require more aggressive policy in the form of longer commitments to zero interest rate policy.

We summarize this discussion with the following proposition.

Proposition 2. *Assume the economy is in the low state. With forward guidance and $\rho = 1$ the expected long-term interest rate is given by*

$$E_t \sum_{T=t}^{\infty} \beta^{T-t} R_{T+1} \quad = \quad \sum_{j=0}^{\infty} \delta(1-\delta)^j \beta^j \frac{\beta^{k_{t+j}}}{1-\beta} \times (\omega^{\pi}_{t|t-1} + r_H).$$

Forward guidance has the following properties:

i. For given beliefs, $\omega^x_{t|t-1}$ and $\omega^\pi_{t|t-1}$, and promises $k_{t+j} = k_{t+i}$ with $j < i$, the promise k_{t+j} produces a larger increase in aggregate demand than the promise k_{t+i}; and

ii. Lower inflation expectations, $\omega^\pi_{t|t-1}$, reduce nominal space and make a given forward guidance policy less effective.

4.2 Results

Figure 3 plots dynamics under the optimal forward guidance policy and learning. As in earlier figures, we plot only the first 20 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, conditional on the low state in period 1, 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. This restraint is delivered by a substantial rise in nominal interest rates after the period of zero interest rate policy.[20]

The fourth panel shows the profile of promises that implements the optimal forward 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. This steady reduction in promises to a commitment of no additional quarters of zero interest rate policy beyond a certain date makes optimal policy well-approximated by a calendar-based promise.

Central to this narrative are inflation expectations, shown in the bottom row of the figure. 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 when the economy returns to the high 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 nominal interest rate relative to the nominal neutral rate. These effects are supported by positive wealth effects from rising output expectations.

4.3 The Insurance Principle

Optimal policy therefore displays an insurance principle. The risk to be insured is a demand shock with uncertain duration. The central bank self insures by making large state-contingent promises for short-duration shocks to ensure inflation expectations don't fall in the event of a long-duration shock. Maintaining the power of monetary policy is the pay-off. 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. Rising inflation and inflation expectations require

[20]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. The appendix explains why alternative assumptions moderate interest rate dynamics but leave our results largely unchanged. We also show our results don't depend on assumptions about the normalizing constants in beliefs or the truncation point used in our numerical solutions.

Figure 3: Optimal Forward Guidance Policy

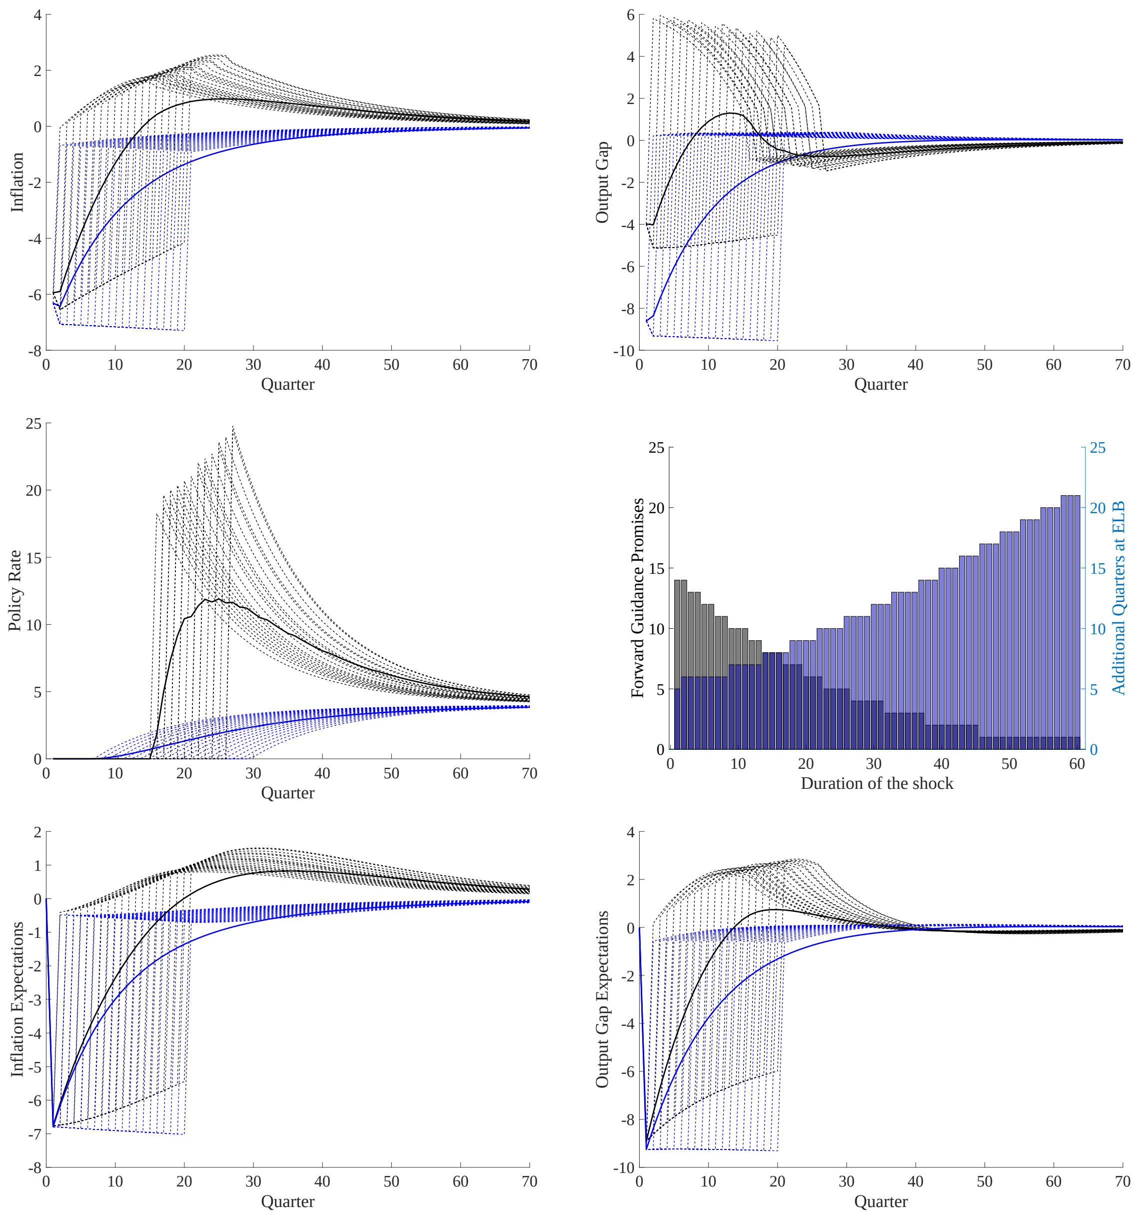

Notes: Optimal forward guidance policy (black) compared to discretion (blue) with learning. Parameter values $g = 0.075$, $\rho = 0.985$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.02$, $\lambda_x = 0.05$, $r_L = -0.003$.

a contraction in aggregate demand after the central bank fulfills the period of promised zero interest rate policy. This is the price of insurance, the insurance premium.

The insurance principle is a unique feature of learning. Households and firms cannot evaluate the general equilibrium implications of forward guidance policy for output and inflation—expectations only adjust slowly to announcements. To prevent deflation and influence demand the central bank must make large forward guidance announcements because of reduced general equilibrium effects of policy. However, over time as households and firms learn about the consequences of policy, the general equilibrium effects of policy on real activity grow. Higher inflation expectations lower real interest rates, stimulating demand through intertemporal substitution. And higher output expectations stimulate demand through standard permanent income wealth effects. The optimal policy manages these delayed but growing effects by tapering promises. Because general equilibrium effects of policy are 'back-loaded', forward guidance policy must be 'front-loaded'.

Learning dynamics therefore shape the profile of forward guidance promises because they shape the economic trade-off that confronts monetary policy. Increasingly persistent beliefs induce larger self-fulfilling deflation and therefore require larger front-loaded stimulus to stabilize the economy. But at the same time, this persistence induces a larger delayed inflation overshooting as the effects of the policy are propagated and amplified. This requires a steeper tapering of forward guidance promises.

That the general equilibrium effects of forward guidance policy are delayed explains why the profile of promises under learning is very different to that under rational expectations. Under rational expectations, because the general equilibrium effects are 'front-loaded', forward guidance promises are 'back-loaded'. Promises to reduce interest rates in future unlikely events have large immediate general equilibrium effects. This means the central bank, initially, needs only to make small promises. Because there are no learning dynamics to amplify and propagate general equilibrium effects, the promises rise with the severity of the recession.

We now characterize these general equilibrium effects of forward guidance policy in a simple analytical example. This renders the above ideas concrete and helps compare our learning model to a rational expectations model and also other recent models of bounded rationality. Section 6 then provides are comparison of the optimal forward guidance policy across a range of models.

5 General Equilibrium Effects of Policy: Simple Analytics

This section provides analytical insight on the insurance principle. In response to forward guidance, we show our model has two properties. First, at the time of announcement, the partial equilibrium effects of policy are substantially larger than the general equilibrium effects. Second, over time the general equilibrium effects accumulate and, depending on beliefs and policy, can become substantially larger than partial equilibrium effects. Strong initial partial equilibrium effects explain the front-loading of policy announcements, while the subsequent strong general equilibrium effects explain the trade-offs.

5.1 A Simple Model

The economy comprises a continuum of households $i \in [0,1]$ with consumption demand functions

$$c_t\left(i\right) = E_t^i \sum_{T=t}^{\infty} \left(m\beta\right)^{T-t} \left[\left(1-\beta\right) x_T - \beta\sigma\left(R_T - \pi_{T+1} - \bar{r}\right)\right]$$

where $0 < m \leq 1$. When $m = 1$ we have the optimal decision rule in the standard New Keynesian model under arbitrary beliefs. When $0 < m < 1$ we have Gabaix's (2020) Behavioral New Keynesian model.[21] The parameter m determines the degree of additional discounting of future outcomes. Goods prices are fixed so that inflation is always zero.[22] This assumption serves simplicity. Allowing price adjustment amplifies the general equilibrium effects of policy that we now discuss, by lowering expected real interest rates even further in response to the policy announcement. Given a monetary policy, the model is closed by the goods market clearing condition

$$x_t = \int_0^1 c_t(i)di.$$

Monetary policy is described as follows. 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.

The economy starts in steady state. Households initially anticipate policy $R_t = \bar{r}$ for all t so that $c_t\left(i\right) = x_t = 0$. Let $c_t^{PE}\left(i\right)$ be the change in consumption induced by implementation of the forward guidance policy described above, holding fixed current income, x_t, and expectations of future income and inflation, $\{E_t^i x_T, E_t^i \pi_T\}_{T=t+1}^{\infty}$. Summing over households gives

$$x_t^{PE} = \int_0^1 c_t^{PE}\left(i\right)di$$

the aggregate partial equilibrium effect of the policy announcement. Using the equilibrium solution for output permits the decomposition

$$x_t = x_t^{GE} + x_t^{PE}$$

[21] This model also provides a representation of other recent models of bounded rationality and zero interest rate policy such as Angeletos and Lian (2018), Bilbiie (2018) and Farhi and Werning (2019).

[22] Formally, we study economic outcomes in the neighborhood of a fixed price equilibrium, so that $\kappa \to 0$.

into general equilibrium and partial equilibrium responses to the policy change.

5.2 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

$$\begin{aligned}\hat{E}_t\left(x_{t+1}|S=L\right) &= (1-\nu)x_L+\nu x_H \\ \hat{E}_t\left(R_{t+1}|S=L\right) &= (1-\nu)\times 0+\nu\bar{r}.\end{aligned}$$

The appendix establishes the following results.

Proposition 3. *Under rational expectations, the forward guidance policy increases equilibrium output with partial and general equilibrium effects*

$$\begin{aligned}x_t^{PE} &= \frac{\sigma\beta}{1-m\beta\left(1-\nu\right)}\bar{r} \\ x_t^{GE} &= \frac{\sigma\left(1-\beta\right)}{\left(1-m\left(1-\nu\right)\right)\left(1-m\beta\left(1-\nu\right)\right)}\bar{r}\end{aligned}$$

for $t\geq 0$. In the limit $\nu\to 0$ the decomposition satisfies

$$\begin{aligned}x_t^{PE} &= \frac{\sigma\beta}{1-m\beta}\bar{r} \\ x_t^{GE} &= \frac{\sigma\left(1-\beta\right)}{\left(1-m\right)\left(1-m\beta\right)}\bar{r}\end{aligned}$$

for $t\geq 0$.

The partial and general equilibrium effects are therefore constant over time. The special case of $m\to 1$ recovers the standard rational expectations aggregate demand function that appears in the canonical new Keynesian model. In this case the general equilibrium effects are unbounded.

Corollary 1. *Under rational expectations, if $m=1$ then*

$$\begin{aligned}\lim_{\nu\to 0} x_t^{PE} &= \frac{\sigma\beta}{1-\beta}\bar{r} \\ \lim_{\nu\to 0} x_t^{GE} &= \lim_{\nu\to 0}\frac{\sigma\left(1-\beta\right)}{\nu\left(1-\beta\left(1-\nu\right)\right)}\bar{r}=\infty.\end{aligned}$$

This is an example of the forward guidance puzzle. The size of the general equilibrium response is proportional to the expected duration of zero interest rate policy, ν^{-1}—the longer the expected duration of zero interest rate policy, the larger the general equilibrium effects. The earlier proposition is Gabaix's (2020) resolution of this puzzle. Higher discounting of anticipated future outcomes moderate general equilibrium effects on output.

5.3 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. Consistent with the numerical results, agents do not know how forward guidance affects the equilibrium value of output at the zero lower bound or their equilibrium values when the economy reverts back to the normal state. Expectations satisfy

$$\begin{aligned} \hat{E}_t\left(x_{t+1}|S=L\right) &= (1-\nu)\omega^x_{t|t-1}+\nu\omega^x_{t|t-1}=\omega^x_{t|t-1} \\ \hat{E}_t\left(R_{t+1}|S=L\right) &= (1-\nu)\times 0+\nu\bar{r}. \end{aligned}$$

We assume that output beliefs are intially at steady state so that $\omega^x_{0|-1}=0$.[23] Output beliefs are revised according to

$$\omega^x_{t+1|t}=\rho(1-\gamma)\omega^x_{t|t-1}+gx_t$$

for $t\geq 0$.[24] This description of beliefs along with the optimal consumption decision rule, and goods market clearing comprise the model. That beliefs are a state variable means that the general equilibrium effects of policy unfold over time.

Proposition 4. *Under boundedly rational responses to policy announcements and learning, the forward guidance policy increases output. The partial equilibrium effect is*

$$x^{PE}_t=\frac{\sigma\beta}{1-\beta\left(1-\nu\right)}\bar{r}$$

for all $t\geq 0$. The general equilibrium effect is

$$x^{GE}_0 = \frac{\sigma\left(1-\beta\right)}{1-\beta\left(1-\nu\right)}\bar{r}$$

[23] Because our analysis of optimal policy implies the constants in beliefs are unaffected by forward guidance we set them here to zero for simplicity.

[24] Recall the Kalman gain satisfies $g=\rho\gamma$.

and

$$x_t^{GE} = \Psi_x x_{t-1}^{GE} + \Psi_r \bar{r}$$

for $t > 0$*, where*

$$\begin{aligned} \Psi_x &= \left[1 - \left(1 - \frac{m(1-\beta)}{1 - m\beta\rho}\right)\gamma\right]\rho \\ \Psi_r &= \frac{(1-\rho(1-\gamma))(1-\beta)}{1 - m\beta(1-\nu)}\sigma. \end{aligned}$$

In the case of the standard model with $m = 1$, the initial partial equilibrium effect is substantially larger than the general equilibrium effect because of the dependence on the household's discount factor. The partial equilibrium effect on impact is identical to rational expectations. The general equilibrium effect on impact is smaller by a factor ν^{-1}, reflecting the assumption of boundedly rational responses to policy announcements. The fact that agents cannot infer the equilibrium implications for future output in response to the policy announcement substantially moderates the stimulatory effect. However, the total general equilibrium effect is governed by a first-order difference equation with the following properties.

Corollary 2. *Under boundedly rational responses to policy announcements and learning, as* $\nu \to 0$ *and* $t \to \infty$ *the following properties hold for the average size of general equilibrium effect:*

1. *If* $m = 1$ *and* $0 \leq \rho < 1$ *then the general equilibrium effect has a finite mean.*
2. *If* $0 < m < 1$ *and* $\rho = 1$ *then the general equilibrium effect has a finite mean equal to the rational expectations equilibrium effect.*
3. *If* $m = 1$ *and* $\rho = 1$ *the general equilibrium effect coincides with rational expectations and is unbounded*

The first result implies the model does not display the forward guidance puzzle. Much like Gabaix, subjective beliefs that are mean reverting imply the future is discounted more heavily than the standard rational expectations model, moderating general equilibrium effects of policy. The second result shows that with persistent beliefs the general equilibrium effect under learning converges to that observed under rational expectations. Even though general equilibrium effects are small on impact, when beliefs are highly persistent they accumulate over time and ultimately converge to the effects predicted by rational expectations. The final result is a special case of the second. With a standard aggregate demand curve and highly persistent beliefs, the general equilibrium effect is unbounded. The forward guidance puzzle appears once again.

What about finite expected durations of zero interest rate policy? Can the general equilibrium effects under learning overshoot the rational expectations effects?

Proposition 5. *If $m = 1$ then there exists parameter values for ρ, γ and ν and a time t^* such that for $t > t^*$ the general equilibrium effect with bounded rationality and learning is larger than that under rational expectations.*

In general, for a given expected duration of zero interest rate policy over-shooting of rational expectations general equilibrium effects are more likely for larger values of the parameters ρ and γ—more poorly anchored expectations. It is this characteristic of learning dynamics that presents stabilization challenges for monetary policy. Self-fulfilling beliefs engender general equilibrium effects which accumulate over time and are potentially large.

To give quantitative content to the result, Figure 4 plots the dynamics of the general equilibrium effect for different assumptions about beliefs. The red dashed line gives the constant general equilibrium effect under rational expectations. The other four lines give different combinations of the belief parameters γ and ρ. Smaller values of each parameter represent better anchored expectations: beliefs are less sensitive to forecast errors and less persistent. Reflecting this, this dotted and solid black lines reveal weaker general equilibrium effects, overshooting the rational expectations effect after five and three years respectively. In contrast high values of both γ and ρ represent more poorly anchored expectations, generating larger general equilibrium effects within one year of the shock.

The optimal forward guidance policy is clearly explained by these properties. Relatively strong partial equilibrium effects on announcement favor aggressive front-loaded stimulus. But because general equilibrium effects depend on the dynamics of expectations and grow over time, subsequent macroeconomic control is more difficult. Indeed, the price of upfront support of the macroeconomy is subsequent restraint in real activity to unwind inflationary pressures from rising long-term inflation expectations. Conversely, the cost of not acting to offset a shock is a growing drag on economic activity through this same force.

5.4 Discussion

These predictions depend on boundedly rational responses to policy announcements that arise under learning. The following situates these implications of learning in the literature.

Level-k reasoning. Our approach follows Preston (2006) and Eusepi and Preston (2010) which analyze central bank communication. It is related to Farhi and Werning (2019). The

Figure 4: General Equilibrium Effects of Forward Guidance

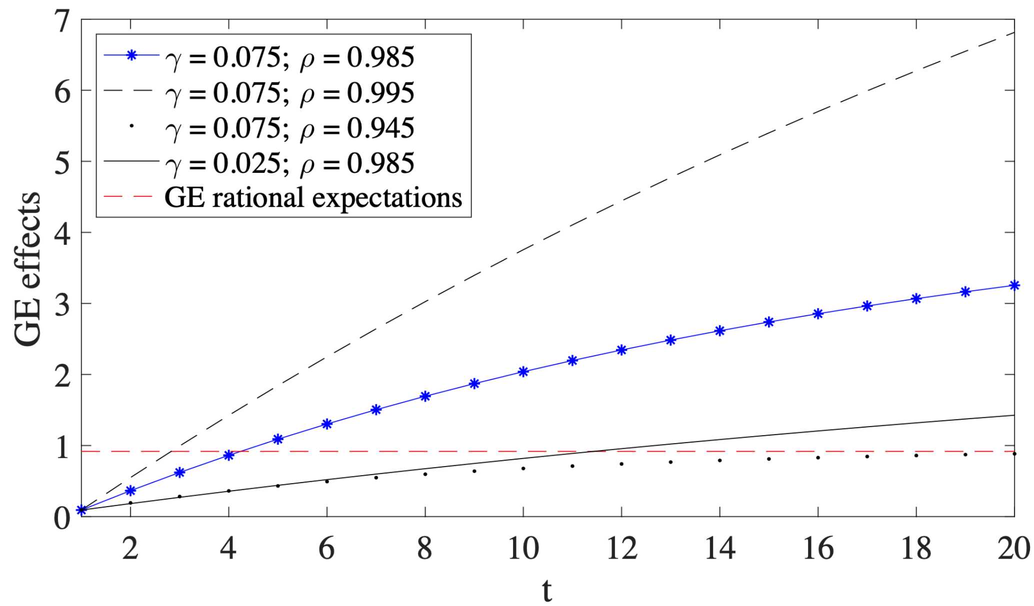

Notes: The figure shows the evolution of general equilibrium effects over time with $m = 1$ and $\nu = 0.1$.

fact that agents don't contemplate how changing future interest rates affect inflation and output gap forecasts, means the impact effects of forward guidance is an example of level-1 reasoning. 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, $k = 1$. The central difference between Farhi and Werning (2019) and us is that beliefs respond to policy over time, generating dynamic general equilibrium effects.

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 rationality and incomplete markets. These models effectively predict higher discounting as in the Gabaix model, so that the partial and general equilibrium effects of forward guidance are almost identical to proposition 1.

6 Applications of the Insurance Principle

This final section explores further the economics of the insurance principle. First, we compare the optimal forward guidance policy with learning to the optimal policy under rational expectations as well as other recent behavioral New Keynesian models. 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

We now characterize the optimal forward guidance policy for Gabaix's (2020) behavioral New Keynesian (BNK) model and Bilbiie's (2018) tractable heterogeneous agent (THANK) model and compare them with the optimal commitment policy under rational expectations and learning. To facilitate comparison, we calibrate the natural rate shock in each model to ensure the same initial decline in output under optimal discretion. For the BNK and THANK models this requires natural rate shocks that are, respectively, 10 times and 4 times larger than the rational expectations and learning models. This is because weaker general equilibrium effects also dampen the economic effects of shocks.[25]

Figure 5 summarizes the findings. The four models are rational expectations (black); BNK model (red); THANK model (turquoise); and the learning model (blue). The top two panels show the dynamics of inflation and output for each of the four models under optimal discretion. The dynamics under rational expectations and learning are familiar. The THANK model has identical dynamics to rational expectations. In contrast, the BNK model has a moderated fall in inflation, reflecting the considerable discounting of future marginal costs when firms set their optimal price in Gabaix's framework.

The middle two panels show the dynamics under optimal forward guidance policy. As in earlier figures, the plotted variables are the central bank's expectations of these variables at the time of the shock, conditional on policy. Because the THANK and BNK models reduce the strength of general equilibrium responses to forward guidance announcements, optimal policy is less successful when compared to rational expectations. The THANK model predicts a small overshooting of inflation, driven by a modest boom in output, giving the appearance of an amplified version of rational expectations dynamics. The BNK model predicts relatively modest stabilization benefits of forward guidance policy.

The bottom panel plots the state-contingent forward guidance promises that generate these paths. Because of increased discounting of the future and reduced sensitivity to real interest rates, forward guidance in the THANK and BNK models is less powerful than the standard rational expectations New Keynesian model. This requires larger promises than under rational expectations.

These results reveal the optimal policy implications of our model are radically different to existing frameworks on two grounds. First, in our model forward guidance policy has much less precise macroeconomic control. Because of imperfect control of the term structure of expectations, forward guidance delivers a strong rebound driven by rising long-term inflation and output expectations. Consequently, it features a tighter monetary policy when exiting the zero lower bound generating a downturn in economic activity. The pronounced overshooting of inflation and subsequent downturn in real activity are not present in other models. Second, learning induces a substantially different profile of forward guidance promises. Our model has large front-loaded promises, because the general equilibrium effects are back-loaded, occurring with a delay. The other models display profiles with back-loaded promises. Even though general equilibrium effects are dampened in the BNK and THANK models, they remain front-loaded, occurring immediately at the time of the announcement

[25] See the online appendix for results which normalize by the size of shock as opposed to the decline of output on impact. Since these models introduce some additional parameters, we take values reported to be the preferred specifications by the authors. In the case of Bilbiie's (2018), his calibration seeks to match the model of McKay, Nakamura, and Steinsson (2016).

Figure 5: Optimal Policy in the Rational Expectations, BNK and THANK Models

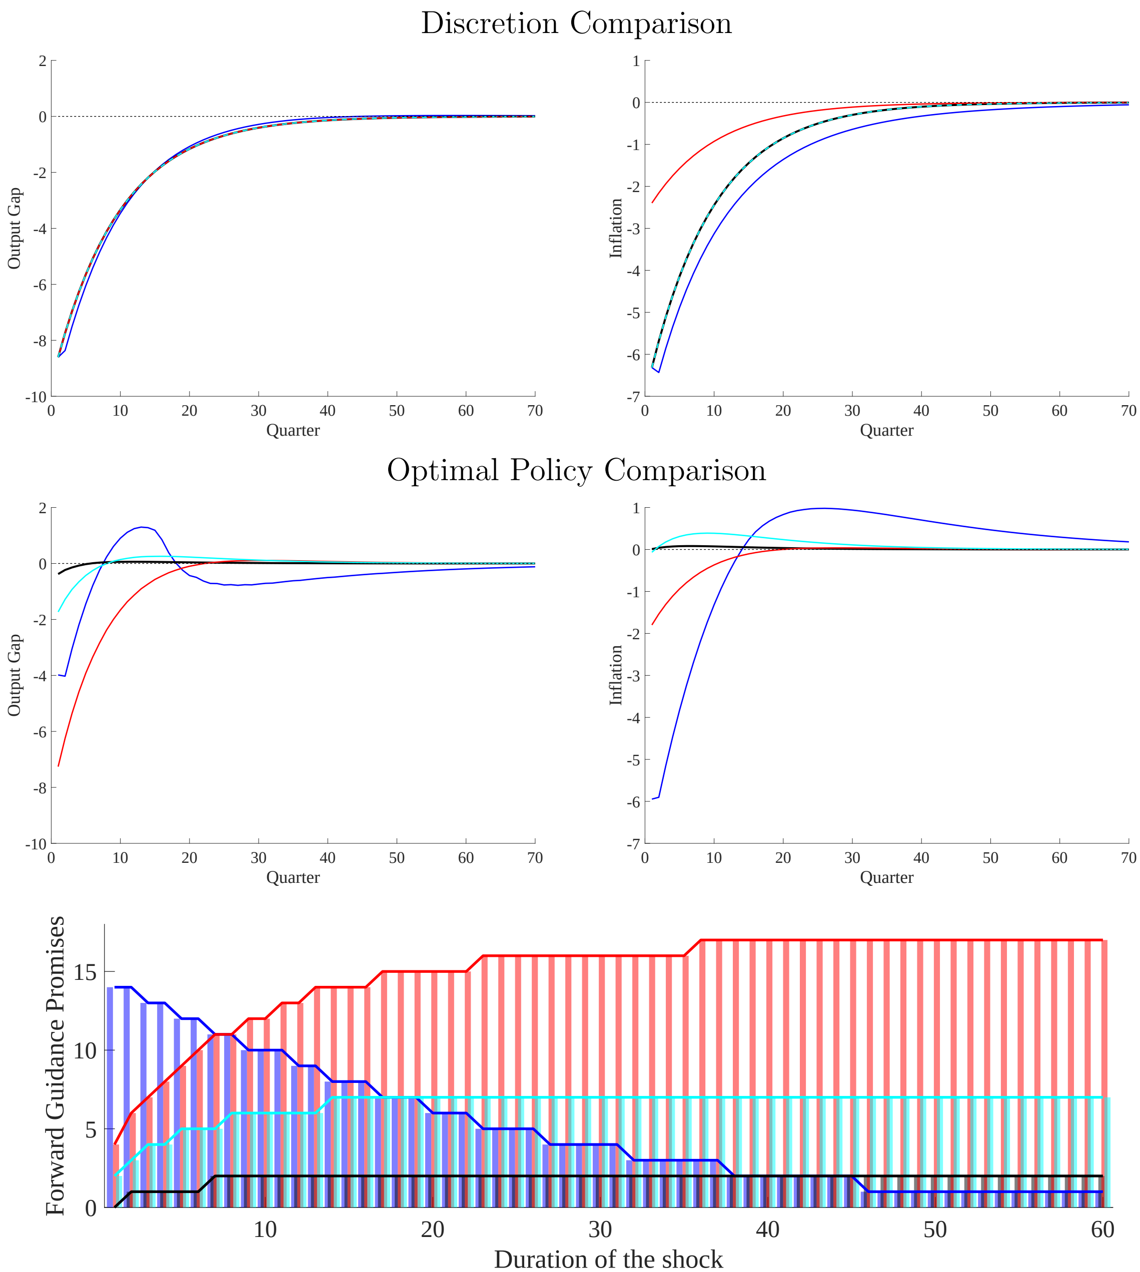

Notes: The natural rate shock is normalized so that it has the same impact in period one on the output gap under discretion for RE (black), learning model (blue), BNK model (red), and a THANK model (cyan). Shared parameter values across all models are $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.02$ and $\lambda_x = 0.05$; RE shock parameter $r_L = -0.003$; learning model parameters $g = 0.075$, $\rho = 0.985$, $r_L = -0.003$; BNK parameters: $M = 0.85$, $M^f = 0.8$, and $r_L = -0.035$. THANK parameters: $M = 0.9701$, $M^f = 1$, $N = 0.843$, and $r_L = -0.012$.

of forward guidance.

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.[26] The path of the policy rate is then optimally determined, consistent with regime 3 in the earlier commitment problem.

Figure 6: Costs and Benefits of Insurance

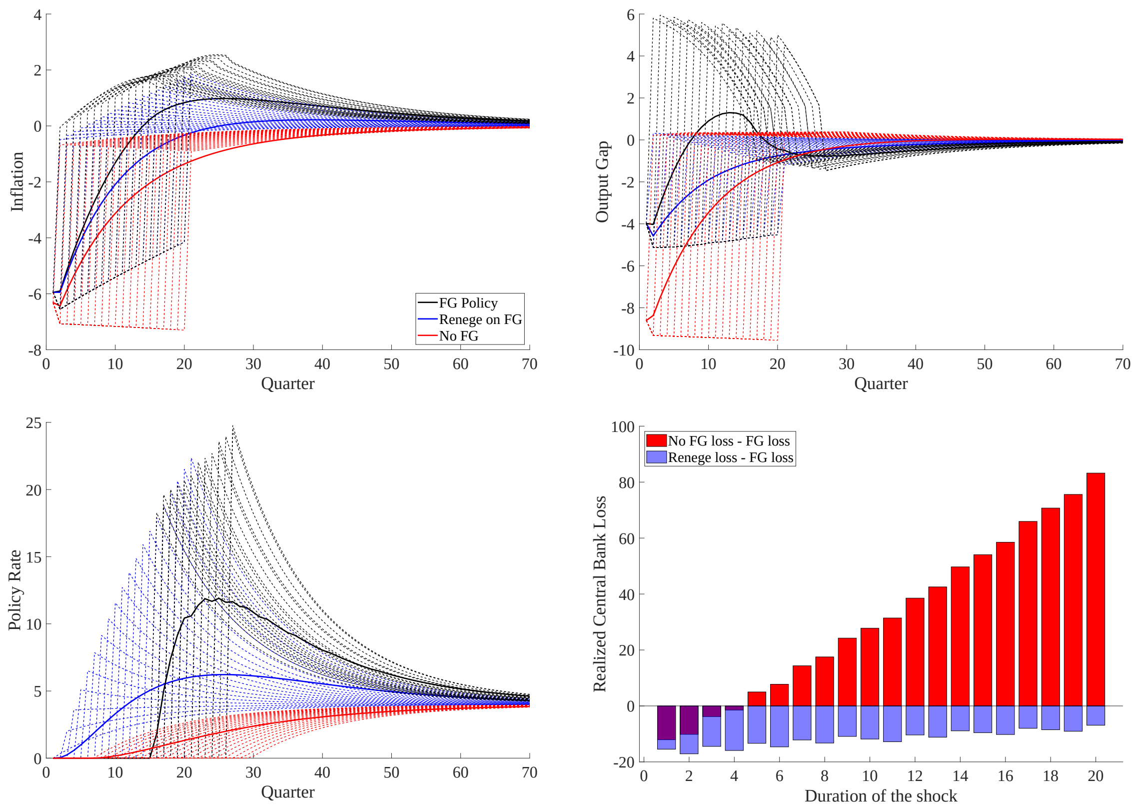

Notes: Parameter values $g = 0.075$, $\rho = 0.985$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.02$, $r_L = -0.003$. Bottom right figure shows a decomposition of the realized benefit of optimal policy. The difference in central bank loss of optimal discretion policy to optimal forward guidance policy is shown in red. The difference in central bank policy of reneging on FG policy to optimal FG policy is shown in blue.

Figure 6 displays the results. The blue lines in each panel show equilibrium outcomes for

[26] 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.

a central bank that reneges on the announced forward guidance. The black and the red lines reproduce earlier results for the optimal forward guidance and no forward guidance policy respectively. As before the solid lines give the expected dynamics conditional on each policy, while the dashed lines denote specific realizations of uncertainty.

Focusing on the solid lines, 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. 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.

The dashed lines for the individual realizations provide additional nuance. As the duration of the shock rises, the output gap at the time of normalization progressively falls, so that longer duration shocks 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 but not yet reneged on.

The bottom right panel decomposes the realized welfare losses conditional on the first 20 possible realizations of uncertainty. The red bars show the welfare loss of the no forward guidance policy relative to the forward guidance policy, the cost of doing nothing. If the shock last four quarters or less, forward guidance policy lowers welfare. For shocks lasting more than four quarters, the welfare losses from not implementing forward guidance rise with the duration of shock. That forward guidance policy can reduce welfare for highly favorable realizations of uncertainty underscores a fundamental trade-off confronting monetary policy. Long-duration shocks are extremely costly when expectations are unanchored. To support expectations in such cases requires immediate and large stimulus, even though favorable realizations of uncertainty make this stimulus costly to unwind—sufficiently costly that it would have been preferable to have not implemented forward guidance policy in the first place.

The blue bars show the value of reneging over and above the value of fulfilling announced commitments. The value of reneging tends to decline (in absolute terms as well as a fraction of the benefit of commitment) as the duration of the shock increases. This reflects two competing forces: the longer the duration of the demand shock, the more optimal policy raises inflation and inflation expectations. This raises the incentive to renege because the policy maker correctly anticipates the subsequent restraint in real economic activity to arrest inflation expectations is larger for longer durations. Balanced against this is that longer duration shocks, which occur with diminishing probability, have smaller state-contingent forward guidance promises. As such they represent a less significant constraint on policy actions, which reduces the value of reneging on those commitments. This second effect tends to dominate.

Viewed this way, the value of reneging is the insurance premium. It represents the cost of committing to forward guidance policy. For shocks lasting 8 quarters or less, this cost is large when compared to the benefit of commitment. Here the insurance premium far exceeds the payoff. 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. And as made clear from this figure, the value to being able to credibly commit to a forward guidance policy yields substantial benefits when faced with persistent shocks.

6.3 'Too little, too late': The costs of delay

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 7 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 inflation expectations and the policy

Figure 7: Optimal Policy: Is it worth the wait?

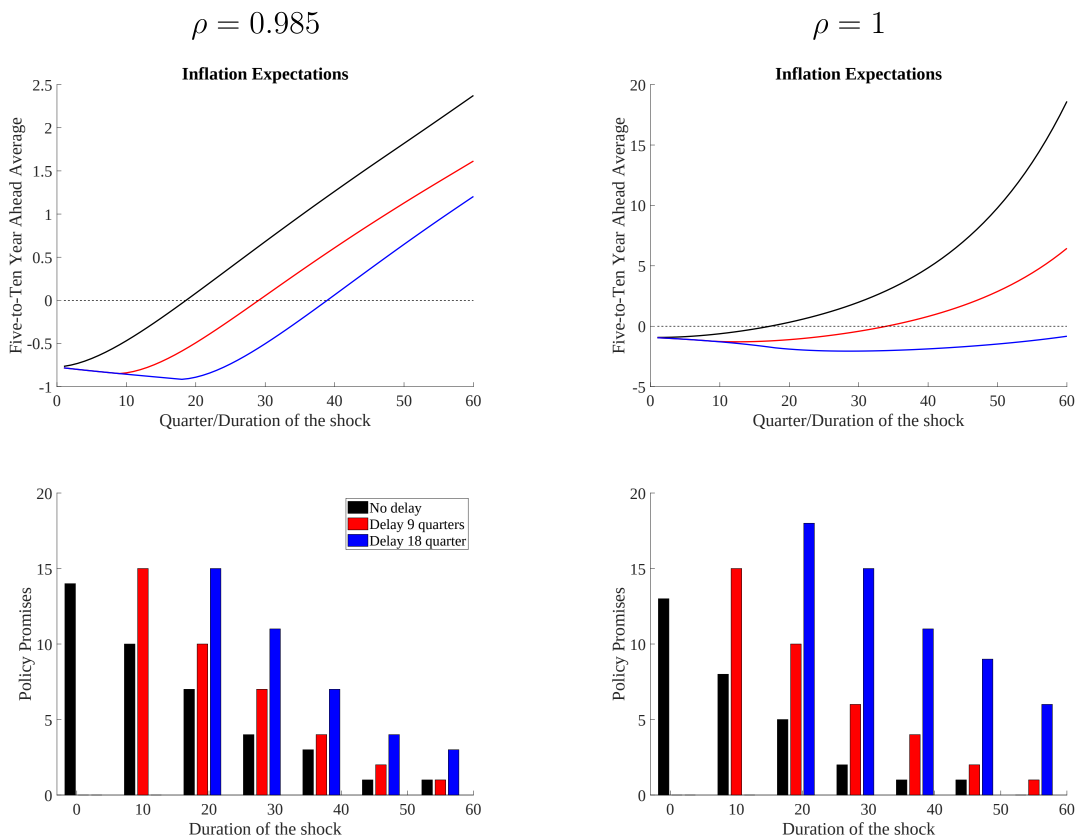

Notes: Shared parameter values $g = 0.075$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\lambda_x = 0.05$, $\kappa = 0.02$, $r_L = -0.003$.

promises. For expectations we plot the average five-to-ten years ahead inflation expectation conditional on a shock lasting to at least the indicated duration.

Start with the dynamics of long-term inflation expectations for the benchmark persistence in beliefs in the left column. Immediate implementation of forward guidance raises long-term inflation expectations in almost linear relation to the duration of shock. In the case of a moderate or long delay in implementation, inflation expectations fall until forward guidance policy is announced and implemented. At the time of implementation, expectations then rise. With highly persistent beliefs, the right column reveals markedly different expectations dynamics. While the impact effects are similar (note the very different scales across the two panels), subsequent dynamics are quite distinct. In the case of delay, persistent beliefs generate self-fulfilling dynamics that lead to on-going falls in long-term inflation expectations. This is most clearly evident in the case of a long implementation delay, where long-term inflation expectations remain below steady state even after 15 years. The same mechanism explains the dramatic rise in long-term inflation expectations in the case of immediate implementation.[27]

Together these first two panels give further perspective on the fundamental challenge posed by unanchored expectations. Highly persistent beliefs propagate the effects of a negative demand shock beyond the period of the crisis. This is most notable in the case of unit root beliefs where drifting beliefs affect inflation expectations into the indefinite future, but applies more generally for sufficiently high persistence. For this reason, policy has less precise control of macroeconomic dynamics. The more persistent are beliefs the more difficult it is to move them; and once they are moving the more difficult it is to restrain them. The general equilibrium effects from belief formation represent a significant constraint on policy.

These properties shape the optimal forward guidance promises shown in the bottom panels. Under the benchmark calibration, the optimal policy profiles show a modest upward shift with increasing implementation delay. Even though the profile of promises doesn't change much, delay is costly because the economy spends a much long time at the zero lower bound. With unit root beliefs, the policy profiles ratchet up sharply with implementation delay, requiring extremely aggressive promises to shift long-term inflation expectations.

Figure 8 shows the associated outcomes for the price level in the case of immediate implementation versus a long implementation delay. To maintain visual clarity, the left column shows dynamics for shocks lasting less than 19 quarters, the right dynamics for shocks lasting 19 quarters or more, conditional on remaining in the low state for at least that long. In the case of delay we assume that if the economy returns to the good state before quarter 19, forward guidance is never implemented. The top row shows the benchmark calibration and the bottom row unit root beliefs. The black lines correspond to an immediate policy response and the blue lines the case of delay. We plot the price level to underscore that reflation is central to good policy, but also to emphasize long-run movements in the price level depend on how well anchored are expectations.

There are two central messages. First, immediate implementation of forward guidance policy reflates the economy. The degree of reflation depends on the severity of the shock and critically the persistence of beliefs. Highly persistent beliefs generate larger increases in the price level because of self-fulfilling expectations—the insurance premium is larger. Second, delay in the implementation of forward guidance creates significant deflationary pressure. This is true even if the crisis ends and forward guidance is never implemented (the left column). The pace of deflation moderates when the natural rate reverts to the high state, but continues until converging to a lower price level, when inflation and output eventually return to steady state. If the shock is long lasting, the eventual implementation of forward

[27] Long-term inflation expectations of 20 percent at first glance might appear extreme. However, the probability of experiencing a large negative demand shock that lasts sixty quarter (15 years!) is 0.002.

Figure 8: Optimal Policy: Is it worth the wait? Price level dynamics

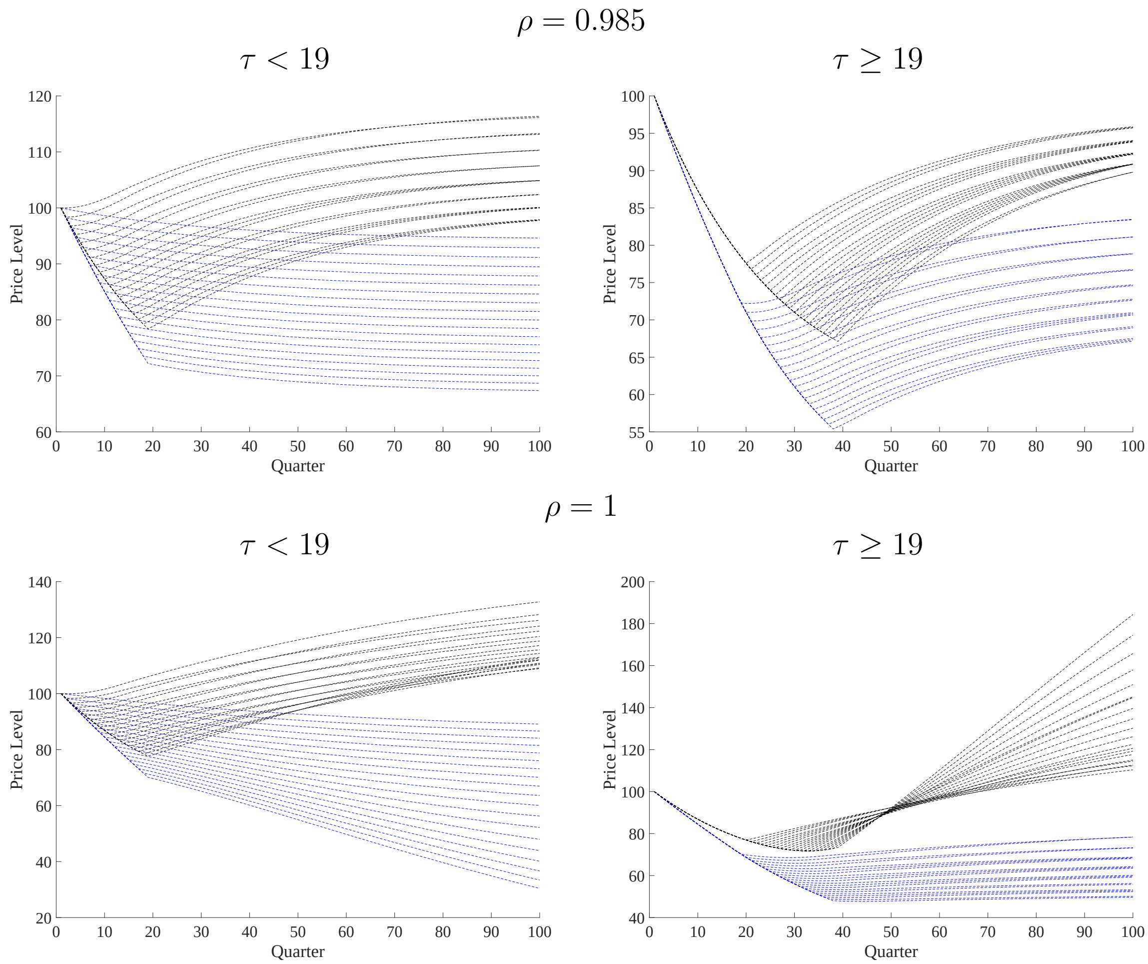

Notes: Parameter values $g = 0.075$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\lambda_x = 0.05$, $\kappa = 0.02$, $r_L = -0.003$.

guidance arrests deflationary pressures, but never generates sufficient inflation to restore the price level to its pre-crisis value.

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., 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.

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.

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.

——— (2016): "The term structure of expectations and bond yields," Staff Reports 775, Federal Reserve Bank of New York.

Crump, R. K., S. Eusepi, E. Moench, and B. Preston (2021): "The Term Structure of Expectations," Staff Reports 992, Federal Reserve Bank of New York.

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.

Evans, G. W., S. Honkapohja, and K. Mitra (2021): "Expectations, Stagnation and Fiscal Policy: a Nonlinear Analysis," Working paper, University of Oregon.

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

Ferrero, G. (2007): "Monetary Policy, Learning and the Speed of Convergence," *Journal of Economic Dynamics and Control*, 39(9), 3006–3041.

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.

Gaspar, V., F. Smets, and D. Vestin (2006): "Adaptive Learning, Persistence and Optimal Monetary Policy," *Journal of the European Economics Association*, 4, 376–385.

Gáti, L. (2021): "Monetary Policy & Anchored Expectations: An Endogenous Gain Learning Model," Working paper.

Gibbs, C., and N. McClung (2020): "Does my model predict a forward guidance puzzle?," *Available at SSRN 3662025*.

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.

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.

Molnar, K., and S. Santoro (2013): "Optimal Monetary Policy When Agents Are Learning," *European Economic Review*, 66, 39–62.

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

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

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

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.