Monetary Policy as Insurance*

Stefano Eusepi† Christopher G. Gibbs‡ Bruce Preston§

Abstract

We study zero interest-rate policy in response to a large negative demand shock of uncertain duration. When individuals must learn about the general equilibrium effects of policy over time the optimal forward guidance policy features an important insurance principle: the optimal policy makes large front-loaded promises to stabilize expectations. Because of this, the optimal policy will appear "too stimulatory" in the event the shock turns out to be transitory—precisely because it is providing insurance against the risk of the shock being persistent. Optimal state-contingent policy is well-approximated by calendar-based forward guidance.

Keywords: Optimal Monetary Policy, Learning Dynamics, Long and variable lags, Forward Guidance

JEL Codes: E32, D83, D84

---

*December 3, 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, NBER Summer Institute Impulse and Propagation Program and the Methods and Applications for Dynamic Equilibrium Models Program, Princeton University, NYU, Columbia University, Duke University, Federal Reserve Board of Governors, Waseda University conference "Financial and Economics Developments: New Challenges and Policy Solutions", National University of Singapore "Macroeconomics Workshop", 2024 Shanghai University of Finance and Economics Macro Workshop, and particularly Hassan Afrouzi, Larry Christiano, Joel Flynn, Cosmin Ilut, Greg Kaplan, Eric Leeper, Karthik Sastry, Ivan Werning, Michael Woodford, our discussant Nigel McClung and especially Chris Edmond for detailed comments. The usual caveat applies. This research was supported by the Australian Research Council under the grant DP210101204.

†Brown University. E-mail: stefano_eusepi@brown.edu

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

§UNSW Business School. E-mail: bruce.preston@unsw.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

When central banks face a crisis of uncertain duration, with current nominal policy rates constrained by the zero lower bound, rational expectations new Keynesian theory prescribes stimulus through state-contingent commitments to future zero interest rate policy at times when policy is no longer constrained. This "lower for longer" policy prescription has the intuitive property that the period of zero interest rate policy to be implemented rises with the length of the crisis. Short-lived crises warrant modest commitments, while progressively longer-lived crises require progressively longer commitments to zero interest rates.

This paper shows that this 'back-loaded' profile of forward guidance commitments will in general be suboptimal. We obtain this result because our model features a departure from standard rational expectations assumptions. Large shocks and new policies make past experience a poor guide to future developments: agents must then learn about the general equilibrium effects of shocks and policy. Combined with a central bank that faces a shock of uncertain duration, optimal forward guidance policy prescribes 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. After a certain duration, optimal policy ceases to commit to further zero interest rate policy, so that calendar-based forward guidance policies well-approximate the fully state-contingent optimal policy.[1]

Front-loaded forward guidance profiles raise output and inflation expectations, making monetary policy more stimulatory in the event of an unfavorable persistent shock. Of course, the optimal policy will appear "too stimulatory" ex post, in the event the shock turns out to be transitory. The policy generates an over-shooting of the inflation target, requiring restraint in real activity to moderate inflation expectations. The central bank optimally accepts poorer stabilization of inflation and output in response to transitory shocks to achieve better stabilization when shocks are more persistent. We call this trade-off the insurance principle.

Central to our results are boundedly rational individuals. Households and firms perfectly understand the implications of forward guidance for the path of future interest rates as well

[1] 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.

as the policy rule used in the post-crisis economy when interest rate policy is normalized. But they must learn about the macroeconomic consequences of the shock and policy response for inflation and output—they cannot evaluate the general equilibrium implications of forward guidance policy. The assumption of bounded rationality is reasonable in situations of the kind we are interested in, when the economy experiences unconventional shocks and policies, about which agents are unlikely to be familiar. The learning problem is not trivial, as forward guidance engenders time-varying non-linear dynamics. Individuals use an approximating statistical model that captures salient features of the true data-generating process, revising beliefs in response to changing economic conditions.

Our results differ from Eggertsson and Wooford's (2003) full-information rational expectations model, which serves as a natural benchmark, in two ways. First, the optimal forward guidance policy requires *front-loaded* and substantially larger state-contingent commitments. Second, policy is less effective, with optimal policy accepting larger fluctuations in inflation and real activity when compared to a rational expectations analysis, which delivers almost complete stabilization of the macroeconomy in response to a negative demand shock. We show these differences reflect distinct general equilibrium effects of forward guidance policy under each belief assumption. Despite having fundamentally different predictions about the design of optimal policy, the transmission mechanism of policy, and short-run stabilization outcomes, they both seek to reflate the economy to a higher long-run price level.

We use the model to decompose inflation and output dynamics in response to forward guidance announcements into three effects:

i. The equilibrium effects without forward guidance policy (which under our assumptions is the rational expectations equilibrium under optimal discretion);

ii. The partial equilibrium effect of forward guidance policy (the change in demand from the announced interest rate path being different to the no forward guidance path, holding fixed current output and expectations of output and inflation); and

iii. The general equilibrium effect of forward guidance policy (the adjustment in output from market clearing and subsequent dynamics from learning).

This decomposition reveals important differences in the short and medium-to-long-run properties of the model. In the short run, on announcement, the general equilibrium effects on output of forward guidance policy are two orders of magnitude smaller than the partial equilibrium effects. Over time, because of learning, persistent general equilibrium effects emerge, leading to long and variable lags of policy of the kind that much concerned Friedman (1960).

These patterns of general equilibrium effect differ markedly to a rational expectations analysis. Because fully rational individuals can evaluate the general equilibrium implications of current and future interest rate policy, at the time of the shock the general equilibrium effects of forward guidance take immediate and large effect—orders of magnitude larger than the partial equilibrium effect, which is identical under both belief assumptions. The central bank exploits these general equilibrium effects by promising larger stimulus only for persistent shocks. Conditional on a transitory shock, where the period of commitment to zero interest rate policy is small, the central bank can easily remove the stimulus, preventing excessive inflation.

Under learning the evolution of general equilibrium effects place constraints on the optimal forward guidance policy. Because general equilibrium effects are initially small, the partial equilibrium channel of policy must assume the burden of economic stabilization policy. And because the partial equilibrium channel is itself modest, mitigating the short-run impacts of the crisis requires policy stimulus initially to be large. But as expectations about economic conditions are revised, the general equilibrium effects become a more dominant source of stimulus. The central bank optimally manages these delayed effects by gradually tapering the size of forward guidance promises for longer-duration shocks. Indeed, profiles of the kind recommended by a rational expectations analysis, that rise with the duration of the crisis, deliver worse short- and long-term stabilization of the economy—activity is too weak in the short run and too strong in the long run, as most stimulus arrives after the expected end of the crisis.

Large-front loaded promises give rise to the insurance principle. The policy ensures stimulatory general equilibrium effects build up gradually to operate in the medium term. These effects provide critical support to the economy in the case of long-duration crisis. But these benefits come with a cost: because the required promises are large, there will be too much stimulus for short-duration shocks.

The results—and specifically the insurance principle—provide a useful framework for understanding recent debate about the efficacy, scale and timing of macroeconomic policy. High inflation during the recovery from a large negative demand shock need not signal inappropriately designed policy—rather it may reflect a central bank that implemented policy to ensure the economy remained resilient in the face of a long-lived crisis. That circumstances proved more favorable than anticipated meant that there was too much stimulus, generating inflation. Ex ante this is optimal. Ex post, inferior stabilization of short-lived shocks is the insurance premium payable to ensure superior stabilization of long-lived shocks. Only if policymakers were certain of a short-lived crisis would smaller scale stimulus be appropriate.[2]

[2] As a specific example, 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

Indeed, we show that in the case of transitory shocks, the central bank would prefer to offer no forward guidance. This same tension provides a rationale for why central banks struggled with their commitments to extended periods of zero interest rate policy: the model reveals the welfare gains from reneging on these commitments and returning to conventional policy are larger for more transitory shocks.

Related literature. Our model displays dampened initial general equilibrium effects from forward guidance announcements because individuals have incomplete knowledge about the aggregate economy. Farhi and Werning (2019), García-Schmidt and Woodford (2019), Woodford and Xie (2019), Gabaix (2020), Angeletos and Lian (2018) and Wiederholt (2015) and Gibbs and McClung (2020) show models of bounded rationality and imperfect common knowledge predict similar attenuation effects.[3] We deviate from these studies in two ways. First, agents must learn about the effects of forward guidance using observed data. Learning engenders delayed general equilibrium effects that grow over time, potentially becoming large, and giving rise to long and variable policy lags. Second, we evaluate the optimal policy response to a demand shock of uncertain duration. The interaction between uncertainty and learning births the insurance principle.

Our general approach to modeling imperfect information and expectations is consistent with a large body of research documenting deviations from the full-information rational expectations. These deviations are well-explained by modeling individuals as statisticians, estimating unobserved economic factors driving the macroeconomy—see for example Evans and Honkapohja (2001), Angeletos and Lian (2023), Crump, Eusepi, Moench, and Preston (2021) and Angeletos, Huo, and Sastry (2020) for reviews of the literature. In contrast to most existing literature, which focuses on short-run expectations and dynamics, our model emphasises learning about the long run: individuals estimate unobserved "trend" components which exhibit high persistence.

Our specific approach to modeling expectations which assumes individuals cannot directly evaluate the general equilibrium effects of changes in the path of interest rates also reflects available empirical evidence. For example, Del Negro, Giannoni, and Patterson (2012) and

the Great Inflation of the 1970s. This paper argues that these views are both essential ingredients of optimal monetary policy.

[3] See Angeletos and Lian (2023) for a discussion about how some of these approaches lead to attenuated general equilibrium effects. An earlier version of our paper computed the optimal forward guidance policy some of these approaches. While general equilibrium effects were diminished, there were not much diminished, leading to similar policy recommendations as Eggertsson and Woodford (2003). For additional evidence on bounded rationality, Farhi and Werning (2019) reviews experimental evidence in games of full information, backing the low level of reasoning that is broadly consistent with our assumptions. Evans, Gibbs, and McGough (2021), motivated by their own experimental evidence, show in a model with both learning and higher-order deductions, that in response to a negative demand shock of the kind considered in this paper agents endogenously coordinate on level-one thinking.

Andrade and Ferroni (2021) show that survey-based inflation and output expectations display weaker responses to forward guidance announcements than what is predicted under rational expectations. A consequence of weaker general equilibrium effects of policy is convergence to rational expectations, if it occurs, is slow, a property of self-referential models emphasized by Marcet and Sargent (1989) and Christiano, Eichenbaum, and Johansen (2024).

We assume that the central bank is fully credible, sourcing reduced effectiveness of forward guidance policy to weaker general equilibrium effects from bounded rationality and learning. This contrasts with Andrade, Gaballo, Mengus, and Mojon (2019) and Bodenstein, Hebden, and Winkler (2019) which emphasize either limited credibility of the central bank, or imperfect information about its preferences. The assumption of full information and perfect central bank credibility is arguably extreme but approximates empirical evidence. Crump, Eusepi, and Moench (2011), Campbell, Evans, Fisher, and Justiniano (2012) and Andrade et al. (2019) show that in response to forward guidance announcements 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.

Finally, our work contributes to the analysis of new Keynesian models with learning in low interest rate environments by Williams (2010), Eusepi (2010), Evans, Guse, and Honkapohja (2008), Evans and Honkapohja (2010) and Evans, Honkapohja, and Mitra (2021). None of these papers study optimal state-contingent forward guidance policy.

2 The Model

We use the simple New Keynesian model of Woodford (2003) and Gali (2008). This facilitates analytical results and comparison to other recent papers on this topic.

2.1 Individual decision rules

With arbitrary beliefs, a log-linear approximation to the optimal individual consumption and pricing decisions of households, $i \in [0,1]$, and firms, $j \in [0,1]$, gives the decision rules

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

$$p_t(j) = \hat{E}_t^j \sum_{T=t}^{\infty} (\xi\beta)^{T-t} \left[(1-\xi\beta) x_T + \xi\beta\pi_{T+1} \right] \quad (2)$$

where x_t denotes the output gap; π_t the inflation rate; R_t the short-term nominal interest rate; and r_t the real natural rate of interest, an exogenous process representing shifts in the propensity to save.[4] This is the only stochastic shock in the model. The operator $\hat{E}_t$ denotes individual firm and household subjective expectations, which can differ from rational expectations.

The first equation is an example of permanent income theory which determines household consumption demand, $c_t(i)$, as the discounted expected value of future income, where the second term captures the effect of variations in the real interest rate on this expected value, because of changes in the nominal interest rate and goods price inflation, and because of exogenous variations in the desired timing of consumption captured by the natural real rate of interest. The parameters $0 < \beta < 1$, $\sigma > 0$ are the household's discount factor and intertemporal elasticity of substitution in consumption. The second equation determines the optimal reset price, $p_t(j)$, as the expected discounted future sequence of output gaps and the inflation rate, which also depends on the exogenous probability $0 < \xi < 1$ of not being able to reset their price in any subsequent period.

2.2 Aggregate implications

Assuming households and firms hold the same expectations, $\hat{E}_t^i = \hat{E}_t^j = \hat{E}_t$ for all i and j, then all households choose the same level of consumption and all firms having the opportunity choose the same price $p_t^* = p_t^*(j)$. Goods market clearing, requiring aggregate consumption demand across households to equal aggregate output,

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

reveals consumption, output and the output gap to be equivalent macroeconomic objects as the natural level of output is constant. Staggered price setting implies the optimal re-set price and inflation are related by $(1-\xi)p_t^* = \xi \pi_t$. Aggregation of the optimal consumption and price decision rules then provides the aggregate demand and supply equations

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

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

where the slope of the aggregate supply curve is $\kappa = (1-\xi\beta)(1-\xi)/\xi$.

[4] See Preston (2005), Woodford (2013) and the Appendix for details.

The structural equations are *unknown* to households and firms, depending as they do on the general equilibrium effects of market clearing and aggregate probability laws. Agents know only their own decision problems, defined by their objectives, constraints and subjective beliefs that we now describe. They do not know that they are identical.

3 The Thought Experiment and Expectations

We are interested in studying economic disturbances that are large, have an uncertain duration and require unconventional policy responses. Each of these elements are relevant to thinking about macroeconomic events in recent decades. Because the shock is large, conventional interest rate policy will be constrained by the effective lower bound on nominal interest rates—unconventional policy is required. Uncertainty about the duration of a crisis matters because forward-looking decision makers will base their plans on expectations about the duration and timing of any stimulus. Policymakers must therefore take this uncertainty into account when designing their policy actions. And by virtue of the economic shock being large, and therefore infrequent, accompanied by innovative policy responses, households and firms will be unfamiliar with the economic context.[5]

The following section develops our approach to modeling uncertainty and imperfect knowledge. We draw a distinction between uncertainty and imperfect knowledge. The central bank faces uncertainty in the form of the duration of the natural rate disturbance but otherwise has complete information about the economic environment. In contrast, while households and firms confront the same uncertainty, they also have imperfect knowledge about the equilibrium implications of the disturbance and concomitant policy response. This information assumption is particularly reasonable in situations of the kind we are interested in, such as the economic consequences of large and infrequent shocks requiring unconventional policy responses, about which agents are likely to be unfamiliar.

3.1 Uncertainty

A simple way to model a shock of uncertain duration is to assume the natural rate evolves according to a two-state Markov process. The natural rate of interest is unexpectedly negative in some period t. In each subsequent period, the natural rate reverts to its steady-state value with a constant probability, $0 < \delta < 1$. When the natural rate is negative, the shock is sufficiently large to ensure that monetary policy is constrained by the zero lower bound on nominal interest rates. For this reason, the central bank will consider forward guidance

[5]Like Lucas (1986) and Sargent (1993) we take the view that rational expectations analysis has "given us theories of dynamics that have their best chance of applying when people are in recurrent situations that they have experienced often before".

policies to lower long-term interest rates. We call this the low state, L, the crisis state. Conditional on being in the low state in any period, the crisis has expected duration $1/\delta$. When the natural rate is at steady state we call this the high state, H. The high post-crisis state is absorbing.

This approach to modeling large shocks of uncertain duration was adopted by Eggertsson and Woodford's (2003) seminal paper on optimal monetary policy at the zero lower bound.[6] Their analysis assumes the central bank to be perfectly credible and households and firms to understand the consequences of forward guidance policy, not only for the path of interest rates, but also inflation and output through general equilibrium effects in every period, into the indefinite future. To isolate the effects of uncertainty on policy design, we maintain the assumption of perfect credibility.

3.2 Information

Agents observe the exogenous natural rate of interest, r_t, and correctly understand its two-state process, including the constant probability of switching. At all times they are aware of the current state of the economy. Agents correctly anticipate that the nominal rate will be set at the zero lower bound while the economy remains at the low state. They understand and incorporate into their decisions state-contingent forward guidance announcements about the path of the nominal interest rate—the monetary authority is perfectly credible. Lastly, individuals know the interest rate rule used to implement conventional policy outside of the zero lower bound regime.

These assumptions are present in a full information rational expectations analysis. We depart from a rational expectations analysis as follows. Because agents do not know the structural equations (3) and (4), which are implications of the aggregation of individual optimal decision rules and market clearing, they have imperfect information about the true statistical model of the aggregate variables they need to forecast to make spending and pricing decisions. In the spirit of Preston (2006) and Eusepi and Preston (2010), agents observe the shock and understand the monetary policy strategy, but because they do not know the structural equations defining the economy, they are unable to evaluate the general equilibrium implications for future inflation and output. Instead agents use an approximating statistical model to learn about the general equilibrium mapping between shocks, policy and economic outcomes.

[6] See also papers by Eggertsson, Egiev, Lin, Platzer, and Riva (2021) and Bilbiie (2019) for further policy analysis in this type of environment.

3.3 Approximating Statistical Model

Before the shock, beliefs are consistent with rational expectations equilibrium in the high state, in which $r_t = r_H = R_t$ and $\pi_t = x_t = 0$. The economy is then subject to a negative demand shock, $r_t = r_L < 0 < r_H$. In each subsequent period the natural rate reverts to the high state r_H with a constant probability $0 < \delta < 1$. Given this information, agents construct forecasts of inflation and output, $z_t = \begin{pmatrix} \pi_t & x_t \end{pmatrix}'$, using the statistical model

$$z_t = \begin{cases} \omega_{0,L} r_L + \omega_t + e_t, & S = L \\ \omega_t + e_t, & S = H \end{cases} \tag{5}$$

and

$$\omega_t = \rho \omega_{t-1} + u_t, \quad 0 \leq \rho \leq 1.$$

The vector $\omega_{0,L}$ captures the effects of the shock to the natural interest rate in the L state, *in absence of forward guidance policy.* This immutable prior is not updated over time and is discussed further below. Our focus in on the vector ω_t, which measures the general equilibrium effects of forward guidance policy that individuals must learn. These effects are time-varying and ultimately mean-reverting. Depending on the parameter ρ, the perceived consequences of policy can be predicted, during the crisis, to persist beyond the duration of the shock to the natural rate. These dynamic effects are approximated by a first-order autoregressive process with innovations u_t. As we discuss below, this simple approximating model captures the salient features of the true data-generating process for output and inflation.

The *i.i.d.* noise disturbance, e_t, captures a perceived "measurement error" so that ω_t is not fully observed: individuals estimate ω_t each period using the observed output gap and inflation. In both states, estimates of the unobserved drift, given information available up to time t, $\omega_{t+1|t}$, are updated each period using the Kalman filter

$$\omega_{t+1|t} = (\rho - g)\, \omega_{t|t-1} + g \left(z_t - \omega_{0,L} r_L \cdot \mathbb{1}_{\{S=L\}} \right), \tag{6}$$

where $\mathbb{1}_{\{S=L\}}$ is an indicator function, taking a value of unity in the low state and zero otherwise. The learning gain $0 < g < 1$, measuring the sensitivity of belief revisions to recent data, is a function of the parameter ρ and individual priors about the covariance of innovations e_t and u_t.[7] The evolution of $\omega_{t+1|t}$ captures how agents learn about the mapping

[7] We assume that the variance-covariances of the innovations obey $\Sigma_u = c \cdot \Sigma_e$, where $c > 0$ is a scalar so that the learning gain is a scalar. See Eusepi, Giannoni and Preston (2024) and the online appendix for a full derivation.

between the announced forward guidance policy and aggregate variables.[8] Forecast surprises lead to revisions in individual forecasts of inflation and output gap.

Because beliefs are consistent with rational expectations equilibrium at the time of the unanticipated shock, so that $\omega_{0|-1} = 0$, the initial forecast error is given by $z_0 - \omega_{0,L} r_L$. Whether the initial forecast error is positive, negative or zero depends on the difference between priors and equilibrium outcomes.

Given an estimate of ω_t, conditional expectations of output and inflation for $T \geq t$ are given by

$$\hat{E}_t z_{T+1} = \begin{cases} (1-\delta)^{T+1-t} \omega_{0,L} r_L + \rho^{T+1-t} \omega_{t|t-1}, & S = L \\ \rho^{T+1-t} \omega_{t|t-1}, & S = H. \end{cases} \tag{7}$$

Recall that δ determines the expected duration of the shock. 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.

3.4 Prior beliefs and no forward guidance benchmark

To fix ideas, consider a benchmark equilibrium with no forward guidance. When $r_t = r_L$ the central bank sets the interest rate equal to zero. When $r_t = r_H$, the central bank resumes conventional interest rate policy. We assume individuals form interest rate projections that are consistent with the optimal "normal times" policy rule so that

$$\hat{E}_t R_{t+j} = r_H + \psi_\pi \hat{E}_t \pi_{t+j} + \psi_x \hat{E}_t x_{t+j} \tag{8}$$

for $j > 0$, where the policy parameters satisfy $\psi_\pi, \psi_x > 0$. We derive this optimal policy later in section 4.

Proposition 1. *Assume no-forward guidance policy. Let $\omega_{0,L} = \bar{\omega}_L \equiv (\bar{\omega}_L^x, \bar{\omega}_L^\pi)'$ be model consistent beliefs. Then,*

i. $\omega_{t|t-1} = 0$ for all t;

ii. the equilibrium perceived and actual law of motion of the economy takes the form

$$z_t = \begin{cases} \bar{\omega}_L r_L < 0 & S = L \\ 0 & S = H, \end{cases}$$

[8] We use the notation $\omega_{t|t-1}$ to emphasize the estimate of $\bar{\omega}_t$ are formed using data available to period $t-1$. This is to avoid simultaneity in determining these estimates and other endogenous variables.

where $\bar{\omega}_L$ is the equilibrium response of output and inflation under rational expectations optimal discretion. The interest rate $R_t = 0$ in the low state and $R_t = r^H$ in the high state.

Proof. Details are in the appendix. The first claim follows from the fact that the forecast error at $t = 0$ is zero in (6). The second claim is immediate from interest rate beliefs being consistent with (8). □

Eggertsson and Woodford (2003) demonstrate this equilibrium can deliver substantial welfare losses. These losses can be reduced by introducing forward guidance policy, which we introduce in the next section. To isolate the effects of forward guidance on learning dynamics, through the paper we assume $\omega_{0,L} = \bar{\omega}_L$. Prior beliefs about output and inflation in the L state in absence of forward guidance policy coincide with rational expectations.

4 Forward Guidance Policy and Equilibrium

In response to the negative demand shock, the central bank's policy strategy comprises two components. The first is state-contingent forward guidance about the duration of zero-interest-rate policy to be implemented after the shock dissipates. The second, occurring after the implementation of these forward guidance commitments, is a return to the optimal pre-crisis conventional policy. We take conventional policy as a fixed feature of the environment.[9]

In this context, we study the optimal forward guidance policy response under imperfect knowledge and learning, and highlight the key differences to a rational expectations economy. We start by studying the structure of economic decisions for an arbitrary forward guidance policy. Assume the central bank can commit to an interest rate policy that is contingent on the duration of the shock. 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 $\tau \in [0, 1, 2, 3, ...]$. These state-contingent promises are assumed to be fully credible and correctly understood by households.

The economic and policy environment induce three regimes. At the time of the negative demand shock, conventional policy becomes constrained by the zero lower bound on nominal interest rates. The central bank announces a state-contingent forward guidance policy and the current interest rate is set equal to zero. This is the first regime. After the demand shock dissipates, with the natural rate rising from the low to the high state, the central bank implements the period of zero interest rate policy that was promised for that specific realization of uncertainty. This is the second regime. After the forward guidance period, the

[9] We could jointly choose the forward guidance and post-crisis policy, but this does not affect the general insights that follow.

central bank optimally chooses conventional interest rate policy given household and firm decisions and beliefs. This is the third regime.

Equilibrium evolution of inflation. We are now in a position to complete the model environment with a discussion of interest rate expectations. Before doing so, use (7) to evaluate output and inflation expectations in (4) to give the aggregate supply curve

$$\begin{aligned} \pi_t &= \bar{\omega}_L^{\pi} r_L \cdot \mathbb{1}_{\{S=L\}} \\ &\quad + \kappa x_t + \frac{(1-\xi)\beta\rho}{1-\xi\beta\rho}\omega^{\pi}_{t|t-1} + \frac{\kappa\xi\beta\rho}{1-\xi\beta\rho}\omega^{x}_{t|t-1} \end{aligned} \tag{9}$$

The top line is the only term dependent on the state. The second term is the contemporaneous relationship between output and inflation. The final two terms capture expectations about future inflation, because of strategic complementarity in price setting, and expectations about future marginal costs which are proportional to output. The more persistent are beliefs, the larger these anticipated effects on pricing decisions and therefore aggregate inflation.

4.1 Regime 1: The Crisis

The first regime occurs when $t < \tau$, during which 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 forward guidance policy to be implemented in all future realizations of uncertainty. We must also describe household expectations about the conduct of interest rate policy after the period of forward guidance.

Long-term interest rate. Given a forward guidance policy $\{k_\tau\}$, the expected path of the nominal interest rate provides the long-term interest rate

$$\hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} R_T = \beta\rho\Psi_t^{\rho}\left(\psi_\pi \omega^{\pi}_{t|t-1} + \psi_x \omega^{x}_{t|t-1}\right) + \beta\Psi_t^{1} r_H \tag{10}$$

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

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

The policy parameters ψ_π and ψ_x are determined by the central bank's optimal "normal times" policy in regime 3. The expression $\Psi_t^{\tilde{\rho}}$ encodes the effect of all future commitments to zero interest rate policy on expected long-term interest rates today.

Each element in the sum reflects one particular contingency that takes the economy to lift-off. For example, consider the case when the economy returns to the high state in period $t+j$ 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 $h = j + k_{t+j}$. At that time, the long-term interest rate is expected to be

$$\hat{E}_{t+h}\sum_{T=t+h}^{\infty}\beta^{T-t}R_T = \frac{1}{1-\beta\rho}\left(\psi_\pi\omega^\pi_{t+h|t+h-1} + \psi_x\omega^x_{t+h|t+h-1}\right) + \frac{1}{1-\beta}r_H.$$

Of course, to obtain the expected interest rate faced in period t, in this particular contingency, we compute the present value, discounting by $\beta^{j+k_{t+j}}$, since one-period interest rates prior to lift-off are equal to zero, and by $\rho^{j+k_{t+j}}$, because beliefs satisfy

$$\hat{E}_t\omega_{t+h|t+h-1} = \rho^{j+k_{t+j}}\omega_{t|t-1}.$$

The actual expected long-term interest rate in period t is then the sum of these interest rates over all possible contingencies.

Aggregate demand. Substituting output and inflation subjective expectations (7), along with interest rate expectations, (10), into (3), provides the aggregate demand equation

$$\begin{aligned} x_t \quad = \quad & \frac{(1-\delta)\left((1-\beta)\bar{\omega}^x_L + \sigma\bar{\omega}^\pi_L\right)}{1-\beta(1-\delta)}r_L + \frac{1-\beta}{1-\beta\rho}\rho\omega^x_{t|t-1} \\ & -\sigma\underbrace{\left(\Psi^\rho_t\beta\rho\left[\psi_\pi\omega^\pi_{t|t-1} + \psi_x\omega^x_{t|t-1}\right] + \beta\Psi^1_t r_H\right)}_{\text{actual nominal rate}} \\ & +\sigma\underbrace{\left(\frac{1}{1-\beta\rho}\rho\omega^\pi_{t|t-1} + \frac{1}{1-\beta(1-\delta)}\left[r_L + \delta\frac{\beta}{1-\beta}r_H\right]\right)}_{\text{neutral nominal rate}}. \end{aligned} \tag{11}$$

The first row captures the general equilibrium effects engendered by prior beliefs about the consequences of the shock, and also the wealth effects from anticipated future income. The second row represents the actual long-term nominal interest rate facing households. The third row measures the neutral long-term nominal interest rate, defined as sum of the long-term natural real rate of interest plus inflation expectations. As always in new Keynesian models, the difference between actual and neutral nominal rates determines the effective stimulus that monetary policy can impart through intertemporal substitution effects, not the level of actual nominal interest rates.[10]

[10]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 because inflation

State-contingent forward guidance therefore lowers the expected long-term interest rate by shifting beliefs about what interest rate will apply in all future realizations of uncertainty. Collecting the terms in r_H we obtain

$$f_t^G = -\sigma\beta\left(\Psi_t^1 - \frac{1}{1-\beta(1-\delta)}\frac{\delta}{1-\beta}\right)r_H, \tag{12}$$

which defines the direct time-t impact effect of stimulus from forward guidance, *holding fixed output and inflation beliefs*. When no forward guidance is offered we have $f_t^G = 0$ for all t but otherwise $f_t^G > 0$.

In general, the larger the promises, the smaller is $\Psi_t^{\tilde{\rho}}$. And 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 effects of forward guidance are non-linear—even though conditional on remaining in the low state, the crisis is always expected to persist for a constant period from that date. This is because the optimal forward guidance policy will in general attach different zero interest rate policy commitments to different contingencies, leading to time-variation in the term $\Psi_t^{\tilde{\rho}}$. Individuals inhabit an environment with structural change from monetary policy, giving rise to a non-trivial learning problem.

The efficacy of zero interest rate policy in lowering the long-term interest rate gap depends on expectations about the policy rate when the economy exits the period of zero interest rate policy. Higher output and inflation expectations imply a steeper path for the expected nominal rate, which curbs economic stimulus. The efficacy of policy also depends on the expected path of the neutral nominal rate. This is because the neutral nominal rate, which depends on inflation and real rate expectations, defines the nominal space available to monetary policy to stimulate economic activity. 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. This is because a longer expected duration of the low state, leads to a larger fall in the long-term real neutral rate.

4.2 Regime 2: The Recovery

The second regime occurs when $\tau \le t \le \tau + k_\tau$, during which the natural rate is r_H but the central bank sets interest rates equal to zero to implement the promised forward guidance

expectations are differenced out in the subtraction of the two nominal quantities.

for that realised contingency. Aggregate demand satisfies

$$x_{\tau+j} = \frac{1-\beta}{1-\beta\rho}\rho\omega^{x}_{t|t-1} - \sigma \underbrace{\left[\frac{\beta^{k_\tau+1-j}}{1-\beta}r_H + \frac{(\beta\rho)^{k_\tau+1-j}}{1-\beta\rho}\left(\psi_\pi\omega^{\pi}_{t|t-1} + \psi_x\omega^{x}_{t|t-1}\right)\right]}_{\text{actual nominal rate}}$$
$$+\sigma \underbrace{\left(\frac{1}{1-\beta}r_H + \frac{1}{1-\beta\rho}\rho\omega^{\pi}_{t|t-1}\right)}_{\text{neutral nominal rate}} \tag{13}$$

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 time when the economy switches back to the high state, and the central bank implements k_τ periods of additional zero interest rate policy. The first term reflects wealth effects from future anticipated income; the second term gives the actual nominal interest rate faced by households; and the third term displays the nominal neutral rate. The difference between the two again determines the effective stimulus from interest rate policy.

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. Again forward guidance imparts non-linearity to the data-generating process. As the economy approaches lift-off the effective stimulus from forward guidance policy is declining. And as before, 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.

4.3 Regime 3: Normal Times

The final regime occurs when $t > \tau + k_\tau$. The negative demand shock has dissipated with the reversion of the natural real rate of interest to its steady state value, and the period of forward guidance ended. The aggregate demand equation is

$$x_t = \frac{1-\beta}{1-\beta\rho}\rho\omega^{x}_{t|t-1} - \sigma \underbrace{\left(R_t + \frac{\beta\rho}{1-\beta\rho}\left[\psi_\pi\omega^{\pi}_{t|t-1} + \psi_x\omega^{x}_{t|t-1}\right] + \frac{\beta r_H}{1-\beta}\right)}_{\text{actual nominal interest rate}}$$
$$+\sigma \underbrace{\left(\frac{1}{1-\beta\rho}\rho\omega^{\pi}_{t|t-1} + \frac{r_H}{1-\beta}\right)}_{\text{neutral nominal interest rate}} \tag{14}$$

with familiar structure. Again, output and inflation expectations are central to nominal space. The critical differences to earlier regimes are that aggregate demand is a linear function of beliefs and current interest rates can be used to manage the expected long-term nominal interest rate.

4.3.1 Optimal policy in normal times.

How does the central bank choose conventional interest rate policy? We assume 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 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.

Monetary policy is therefore chosen to minimize this loss subject to the constraints implied by private behavior: that is, the belief updating equation (6), the aggregate supply constraint (9) and the aggregate demand constraint (14). This is a standard linear-quadratic policy problem, but for the consistency requirement on beliefs, which requires solving for a fixed point. The appendix reports the first-order conditions from which the following result follows.

Proposition 2. *Assume a stationary solution exists. The optimal policy problem implies interest rates are set according to the policy function*

$$R_t \quad = \quad r_H + \psi_\pi \omega^\pi_{t|t-1} + \psi_x \omega^x_{t|t-1} \tag{16}$$

where the policy parameters

$$\begin{aligned} \psi_\pi &= \psi_\pi(\sigma, \xi, \lambda_x, g, \beta, \rho) \\ \psi_x &= \psi_x(\sigma, \xi, \lambda_x, g, \beta, \rho) \end{aligned}$$

represent a fixed point of the policy problem and are therefore functions of underlying primitives.

Proof. A sketch of the proof proceeds as follows. Because the constraints contain only predetermined variables, the Lagrange multipliers on these constraints must be non-predetermined variables. The first-order conditions can then be solved such that all variables are linear functions of the state variables $\omega^\pi_{t|t-1}$ and $\omega^x_{t|t-1}$. In particular, the interest rate solution takes the form

$$R_t = r_H + \tilde{\psi}_\pi(\sigma, \xi, \lambda_x, g, \beta, \rho; \psi_x, \psi_\pi)\omega^\pi_{t|t-1} + \tilde{\psi}_x(\sigma, \xi, \lambda_x, g, \beta, \rho; \psi_x, \psi_\pi)\omega^x_{t|t-1}.$$

The central bank then chooses $\{\psi_\pi, \psi_x\}$ to satisfy the restrictions $\psi_\pi = \tilde{\psi}_\pi(\sigma, \xi, \lambda_x, g, \beta, \rho; \psi_x, \psi_\pi)$ and $\psi_x = \tilde{\psi}_x(\sigma, \xi, \lambda_x, g, \beta, \rho; \psi_x, \psi_\pi)$. This equilibrium can be implemented by the rule (16). □

Optimal conventional policy predicts the interest rate to be a function of the level of the natural interest rate and individual beliefs about the output gap and inflation.

4.4 Perceived and Actual Law of Motion

Inspecting equations (9) and (11)-(14), we can express the data-generating process for the output gap and inflation as

$$z_t = \begin{cases} \bar{\omega}_L r^L + \mathcal{T}_t\left(\omega_{t|t-1}\right) & S = L \\ \mathcal{T}_t\left(\omega_{t|t-1}\right), & S = H, \end{cases}$$

taking the same form as the individual's statistical model. The time-varying function $\mathcal{T}_t(\cdot)$ reflects two model properties. First, shifting expectations affect the economic relationships that agents are attempting to approximate using their statistical model—an example of what Marcet and Sargent (1989) call a self-referential economy. Second, forward-guidance policy on announcement and implementation has time-dependent effects on aggregate output and inflation.

The individuals' forecasting model is only an approximation to the true mapping between shocks, policy and aggregate outcomes across regimes. Subjective beliefs $\omega_{t|t-1}$ generally do not coincide with the objective probability distribution $\mathcal{T}_t$. In particular, the law of motion of $\mathcal{T}_t$ is regime-dependent and non-linear, more complex than a first-order autoregressive process.[11] The following result sheds light on the 'distance' between perceived and actual law of motion, and the convergence properties of the learning process.

Corollary 1. *Let $t = \tilde{t}$ denote the date when the economy switches to regime 3. For parameter values and forward guidance promises ensuring $\tilde{t}$ is finite and that π_t and x_t are bounded at $t = \tilde{t}$, then*

1. *the economy converges to rational expectations equilibrium where $R_t = r_H, x_t = \pi_t = 0$ and*

$$\lim_{t\to\infty} \mathcal{T}_t\left(\omega_{t|t-1}\right) - \omega_{t|t-1} = 0;$$

[11] See the Appendix for details.

2. at any time $t \geq 0$ we have

$$\lim_{T\to\infty} E_t \mathcal{T}_T \;=\; \lim_{T\to\infty} \hat{E}_t \omega_T = 0,$$

where E_t denotes model-consistent expectations.

The belief that $\lim_{T\to\infty} \hat{E}_t\omega_T = \rho^{T-t}\omega_{t|t-1} \to 0$ is consistent with the true data-generating process: the effects of the shock and forward guidance policies are transitory and beliefs converge to rational expectations. The condition required for this convergence result is that, given the expected duration of the crisis, the central bank has sufficient nominal space to support the economy's return to steady state in the aftermath of the shock. The amount of stimulus required depends on the size and persistence of the shock, the parameters that shape beliefs, and the post-crisis policy rule. These conditions are met for a large set of parameters commonly used in the literature.[12]

5 General Equilibrium Effects of Forward Guidance Policy

Rational expectations equilibrium analysis predicts large general equilibrium effects from policy. In Eggertsson and Woodford (2003), modest promises about the future conduct of policy give rise to large effects on current output and inflation. Households and firms anticipate that commitments to keep interest rates low in the future, when the central bank is no longer constrained by the zero lower bound, generate higher equilibrium inflation and output at that time. These future effects are telescoped to the present through higher consumption demand, sustained by expectations of lower real interest rates and higher permanent income, and higher desired prices, because of strategic complementarity in price setting.[13]

We now show that our model with learning dynamics gives rise to fundamentally different general equilibrium effects of policy, radically altering the transmission mechanism of forward guidance announcements.

General equilibrium effects at time 0. At the time of the shock, $t = 0$, aggregating individual consumption demand (1), under maintained assumptions, provides aggregate consumption demand in the low state

$$c_0 \;=\; (1-\beta)x_0 + \beta\bar{\omega}^x_L r_L + \beta f^G_0, \tag{17}$$

[12]In our simulations, the condition was only ever violated when the shock is highly persistent, initial beliefs are optimistic relative to the actual impact effect of the shock, beliefs have a high constant gain and persistence, no forward guidance is offered, and the post-crisis policy rule fails to satisfy the Taylor principle.

[13]Farhi and Werning (2019) give formal expression and quantification of these ideas, in a discussion of the forward guidance puzzle, demonstrating that the general equilibrium effects can be two orders of magnitude larger than the partial equilibrium effect of monetary policy.

where, recalling equation (12), the term βf_t^G defines the exogenous *partial equilibrium* effect of the forward guidance announcement relative to a policy of no forward guidance: the direct effect on consumption, holding fixed income and interest rate expectations after the period of zero interest rates.[14] We define the *general equilibrium* effects of forward guidance policy as

$$x_0^{GE} \quad = \quad x_0 - \bar{\omega}_L^x r_L - \beta f_0^G = (1-\beta) f_0^G.$$

The first equality indicates we exclude $\bar{\omega}_L^x r_L$ from our definition, since it measures the effects of the shock in absence of forward guidance policy. The second equality obtains after using (17) and the goods market clearing condition, $c_0 = x_0$. The initial general equilibrium effect of forward guidance policy is smaller than the partial equilibrium effect by a factor of $1 - \beta$. [15] Using the aggregate supply curve, (9), the general equilibrium effect on inflation is

$$\pi_0^{GE} \quad = \quad \kappa f_0^G$$

with magnitude tied to the slope of the Phillips curve κ.

If there were no learning dynamics, then the model predicts the general and partial equilibrium effects of forward guidance to be constant at the above values. The model would then be a special case of Farhi and Werning's (2019) analysis of level-k thinking, in which households engage in only level-one deductions. We now show that learning generates several new implications. Specifically, the general equilibrium effects of policy are inherently dynamic and endogenous, depending on belief revisions and the policy regime. In period $t = 0$, agents use (6) to obtain their initial estimate of the impact of the forward guidance policy on the aggregate economy, which provides the estimate $\omega_{1|0} = \left(g f_0^G, g\kappa f_0^G\right)' > 0$. They then gradually incorporate the general equilibrium effects of the forward guidance announcement.

Dynamic general equilibrium effects. Output and inflation can be decomposed as

$$x_t \quad = \quad \bar{\omega}_{0,L}^x r_L + f_t^G + x_t^{GE}$$

$$\pi_t \quad = \quad \bar{\omega}_{0,L}^\pi r_L + \pi_t^{GE}.$$

The first time-invariant component determines the equilibrium output and inflation that would emerge in absence of forward guidance announcements. As shown above, this equi-

[14] Recall at $t = 0$ we have $\omega_{0|-1} = 0$, so income and interest rate expectations are constant.

[15] In absence of forward guidance policy, the general equilibrium effect is zero.

librium implies $f_t^G = x_t^{GE} = \pi_t^{GE} = 0$ for all t. The second component of output is the partial equilibrium effect of the forward guidance announcement. The final component for both variables is the dynamic general equilibrium effect from belief updating,

$$
\begin{aligned}
x_t^{GE} &= (1-\beta) f_t^G + \chi_t^x \omega_{t|t-1}^x + \chi_t^\pi \omega_{t|t-1}^\pi \\
\pi_t^{GE} &= \kappa \left(f_t^G + x_t^{GE} \right) + \frac{(1-\xi)\beta\rho}{1-\xi\beta\rho} \omega_{t|t-1}^\pi + \frac{\kappa\xi\beta\rho}{1-\xi\beta\rho} \omega_{t|t-1}^x,
\end{aligned}
\tag{18}
$$

where

$$
\begin{aligned}
\chi_t^x &= \left[\frac{1-\beta}{1-\beta\rho} - \sigma\beta\Psi_t^\rho \psi_x \right] \rho, \\
\chi_t^\pi &= \left[\frac{\sigma}{1-\beta\rho} - \sigma\beta\Psi_t^\rho \psi_\pi \right] \rho.
\end{aligned}
$$

The latter two time-varying parameters govern the link between output and inflation beliefs and the dynamic general equilibrium effects. They depend on the persistence of expectations, ρ, and the policy response. Longer promises of future zero interest rate policy increase the coefficients, while tighter expected policy on lift-off, a higher ψ_x or ψ_π, reduce it. Inflation responds contemporaneously to a change in output, making the slope the aggregate supply curve an important determinant of overall size of general equilibrium effect of policy.

From the Kalman filter updating equation, (6), forecast revisions evolve according to

$$
\begin{aligned}
\omega_{t+1|t}^x &= (\rho - g)\, \omega_{t|t-1}^x + g \left(x_t^{GE} + f_t^G \right) \\
\omega_{t+1|t}^\pi &= (\rho - g)\, \omega_{t|t-1}^\pi + g \pi_t^{GE}.
\end{aligned}
\tag{19}
$$

Forecast surprises are driven by both partial and general equilibrium effects of the forward guidance stimulus. Together (18) and (19) comprise a non-linear system of equations that characterise the evolution of general equilibrium effects in the low state of the economy.[16] Aggregate effects of the forward guidance announcement lead to revised beliefs. These in turn produce general equilibrium effects.

Simple Economy. To provide a pelucid analytical characterization of these general equilibrium effects we simplify further. The next section shows the general equilibrium effects in the full model. First, assume that the central bank promises a fixed period of forward guidance, regardless of the length of recession. Call this k, so that $\chi_t^x = \chi^x$. The

[16] The appendix discusses the general equilibrium effects in the remaining regimes.

expression that encodes the forward guidance commitments for each contingency becomes the constant

$$f_t^G = f^G = \frac{1-\beta^k}{1-\beta}\frac{\delta\sigma\beta}{1-\beta(1-\delta)}r_H.$$

Under this policy, conditional on being in the low state, dynamics are given by a set of linear equations. Second, assume prices are fixed, so that $\pi_t = \omega^{\pi}_{t|t-1} = 0$. The following proposition describes the dynamic general equilibrium effects on the output gap in response to the forward guidance announcement.

Proposition 3. *Under maintained assumptions, the dynamic general equilibrium effect of forward guidance is given by the autoregressive process*

$$x_0^{GE} = (1-\beta)f_G; \quad x_1^{GE} = x_0^{GE} + \chi^x g f^G$$

and, for $t > 1$

$$x_t^{GE} = (1-\hat{\rho})\bar{x}_L^{GE} + \hat{\rho}x_{t-1}^{GE}$$

where $\hat{\rho} \equiv \chi^x g + \rho - g$ *and*

$$\bar{x}_L^{GE} \equiv \left(1-\beta+\frac{2-\beta}{1-\hat{\rho}}g\chi^x\right)f^G.$$

The parameters that regulate the updating of beliefs shape the dynamic behavior of general equilibrium effects induced by forward guidance promises. If beliefs are unresponsive to forecast errors, so that $g = 0$ and no learning occurs, then the general equilibrium effects of forward guidance policy are constant over time. The equilibrium outcomes then correspond to *a level-1 solution* of the model in which $x_t^{GE} = \bar{x}^{GE} = (1-\beta)f^G$. For $g > 0$, forward guidance policy produces a positive initial forecast error, inducing general equilibrium effects x_t^{GE} that are *increasing over time*, provided $\chi^x > 0$. Their magnitude increases with the duration of the crisis, δ, with forward guidance promises, k, and decreases with policy aggressiveness after the period of zero interest rate policy, ψ_x.

Together with χ^x, the size and persistence of general equilibrium effects depend on the persistence of beliefs, ρ, and the learning gain, g. Depending on parameter values x_t^{GE} can grow without bound, conditional on the economy remaining in the low state.[17] The simple example llustrates two key model features. First, general equilibrium effects of stimulus are delayed. Second, they can grow substantially with the duration of the crisis.

Inflation dynamics. While more complex, general equilibrium effects with flexible

[17]The process converges to a finite mean when $0 < g\chi^k < 1-\rho+g$. When beliefs have a unit root, $\rho = 1$, *and* either $k \to \infty$ or $\psi_x = 0$, we have $\hat{\rho} = 1$.

prices share similar, but amplified, behavior. Inspecting equation (18), an increase in inflation expectations, through a rise in $\omega^{\pi}_{t|t-1}$, reinforces the general equilibrium effects of policy by lowering the expected real rate of interest, when $\chi^{\pi}_t > 0$. Also, higher output and inflation expectations increases inflation through the aggregate supply curve, further raising inflation expectations. Finally, from (19), forward guidance announcements affect inflation and inflation expectations only though general equilibrium effects.

6 Optimal Policy Response: Front-loaded Stimulus

This section presents the main results of the paper. We start by demonstrating the optimal forward guidance policy is 'front-loaded'. The central bank commits to large stimulus for crises with short duration, with progressively smaller commitments to zero interest rate policy as the duration of the crisis lengthens. This policy prescription contrasts with the predictions of a rational expectations analysis in which progressively longer periods of zero interest rate policy are associated with longer-duration crises. We show these differences stem from distinct general equilibrium properties of monetary policy under each belief assumption, and that optimal front-loaded stimulus ensures monetary policy provides a form of macroeconomic insurance, a property that is absent under perfect information.

6.1 The Optimal Forward Guidance Policy

Taking as given the process for the natural real rate, the central bank chooses the forward guidance policy $\{k_\tau\}$ to minimize the loss (15) such that the aggregate supply curve (9), belief updating equations (6) are always satisfied and

1. When $t < \tau$ and $r_S = r_L$, the aggregate demand curve (11);
2. When $\tau \leq t \leq \tau + k_\tau$ and $r_S = r_H$, the aggregate demand curve (13); and
3. When $\tau + k_\tau < t$ and $r_S = r_H$, the aggregate demand curve (14), and policy rule (16)

are all satisfied.

Because the problem is non-linear, we use a numerical procedure to determine the optimal policy. Details can be found in the appendix.

Calibration. 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 to the stead-state value of $r_H > 0$ with probability 0.1 in each period. This implies that at any point in time, conditional on being in the low state, the crisis is expected to continue for 10 quarters. The steady-state value of the natural real rate is assumed to be 4 percent per annum, so that $\beta = 0.99$. Agents' perceived persistence of the unobserved process ω_t is set to $\rho = 0.95$. This implies a half-life of just over three years (thirteen quarters), reflecting the expectations that forward guidance policy has an impact beyond the expected duration of the shock. The learning gain, regulating the size of revisions to expectations, is $\bar{g} = 0.15$, capturing substantial uncertainty about ω_t (i.e. a low signal-to-noise ratio). The remaining parameters are set to values similar to Eggertsson and Woodford (2003) to facilitate comparison with earlier results in the literature. We set $\sigma = 0.5$, $\kappa = 0.02$, $\lambda_x = 0.05$.

6.2 'Front-Loaded' Stimulus

Figure 1 plots dynamics under the optimal forward guidance policy and learning. The blue thin lines show the trajectory of the economy for each realization of uncertainty—they are ex post impulse response functions for a shock of a particular duration. In each panel, the solid blue lines give the expected trajectory of each variable, at the time of the shock and conditional on being the low state r_L.[18] This provides a useful ex ante summary of the general character of economic outcomes from the perspective of the policy maker. The black line shows the expected equilibrium outcomes with no forward guidance policy.

The first panel shows the profile of promises that implements the optimal 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 as the duration of the shock increases, leads to a commitment of no additional quarters of zero interest rate policy beyond a certain date. This makes optimal policy well-approximated by a calendar-based forward guidance policy.[19]

The remaining panels show the dynamics of output, inflation, and the policy rate, during and after the zero lower bound period. The expected paths of output and inflation reveal forward guidance stimulus moderates the effect of a large negative demand shock. This comes at the cost of overshooting, delivering a persistent increase in inflation above the long-run target, and a downturn in output in response to tighter policy on exit from the zero interest rate policy. On average, interest rates are expected to remain at zero for a period that well-exceeds the expected duration of the crisis of ten quarters. On normalization, when the natural rate returns to the high state, interest rates rise sharply to deliver the required restraint of real activity to moderate inflation pressure. The interest rate also overshoots its long-run equilibrium.

[18] That is, the average is calculated by weighting each impulse response function by the probability of occurrence and summing.

[19] Under benchmark assumptions, the optimal calendar-based policy commits to 18 periods of zero interest rate policy, delivering 97.5 percent of the welfare gain of the fully optimal policy.

Figure 1: Baseline Optimal Policy

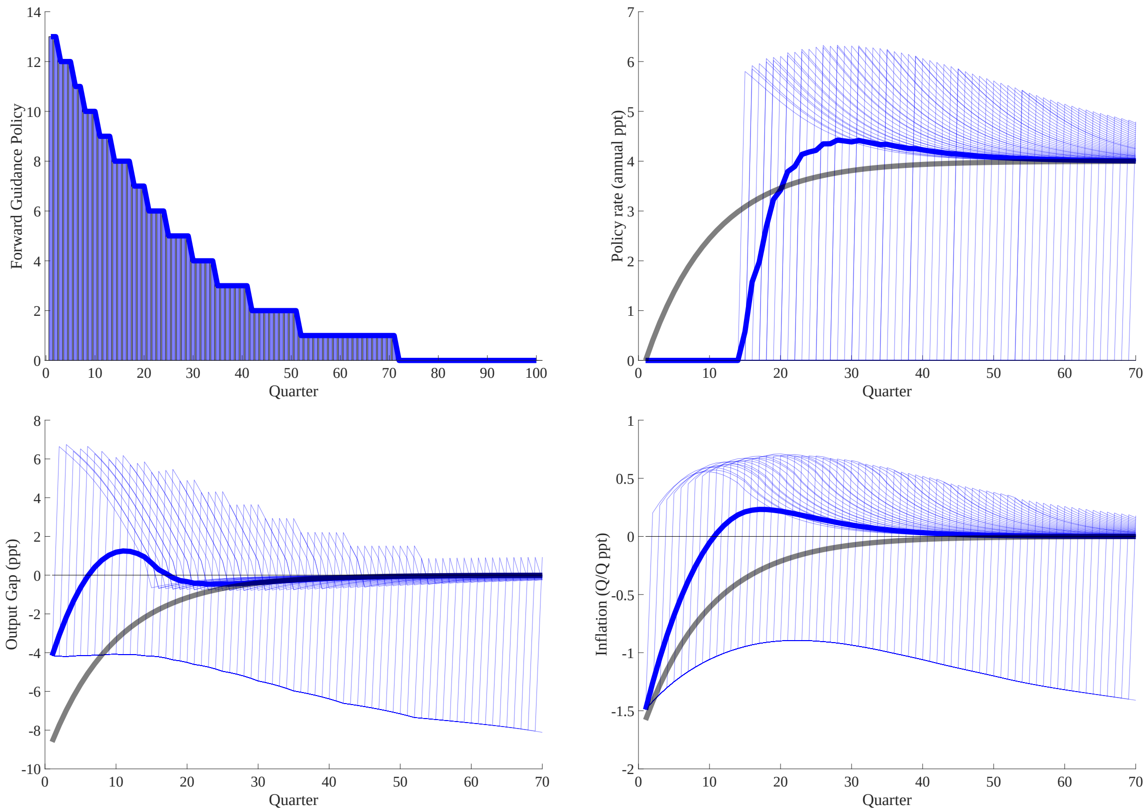

Notes: Optimal forward guidance policy (blue) compared to discretion (black) under RE or learning. Parameter values $g = 0.15$, $\rho = 0.95$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.02$, $\lambda_x = 0.05$, $r_L = -0.003$.

The state-contingent impulse response functions reveal three additional properties of optimal policy. First, when the economy reverts to the high state, output experiences a sizable boom. The expansion is largest for a shock of short duration and gradually shrinks with the duration of the shock. Second, the response of inflation is delayed: on reaching the high state inflation is initially subdued but gradually increases. Even with the progressive tapering of forward guidance stimulus, crises of very long duration are associated with substantial inflation pressure, despite the considerably weaker real economic activity on return to the high state. The peak of inflation is reached for intermediate durations of crises, as inflation in the low state steadily increases. Third, for longer durations, while in the low state both inflation and output approach the equilibrium under no-forward guidance.

The interplay between the delayed effects of policy and the uncertainty surrounding the duration of the crises induces long and variable lags in the policy response, creating additional trade-offs for the policymaker. These features shape the optimal forward guidance profile under learning.

6.3 General Equilibrium effects: The Long and variable lags of policy

The top two panels of Figure 2 decompose the expected dynamics of output and inflation into the three components identified in section 5: The equilibrium response to the shock with no forward guidance (light blue bars), the partial (red bars) and general (dark blue bars) equilibrium effects of the optimal forward guidance policy.

Consistent with the simple example, the general equilibrium effects of forward guidance policy are initially very weak, forcing the partial equilibrium channel of policy to assume the burden of stabilizing output and inflation at the time of the shock. And because this partial equilibrium channel is itself modest in magnitude, the central bank requires large forward guidance commitments.

Figure 2: PE vs GE Effects
Optimal Policy

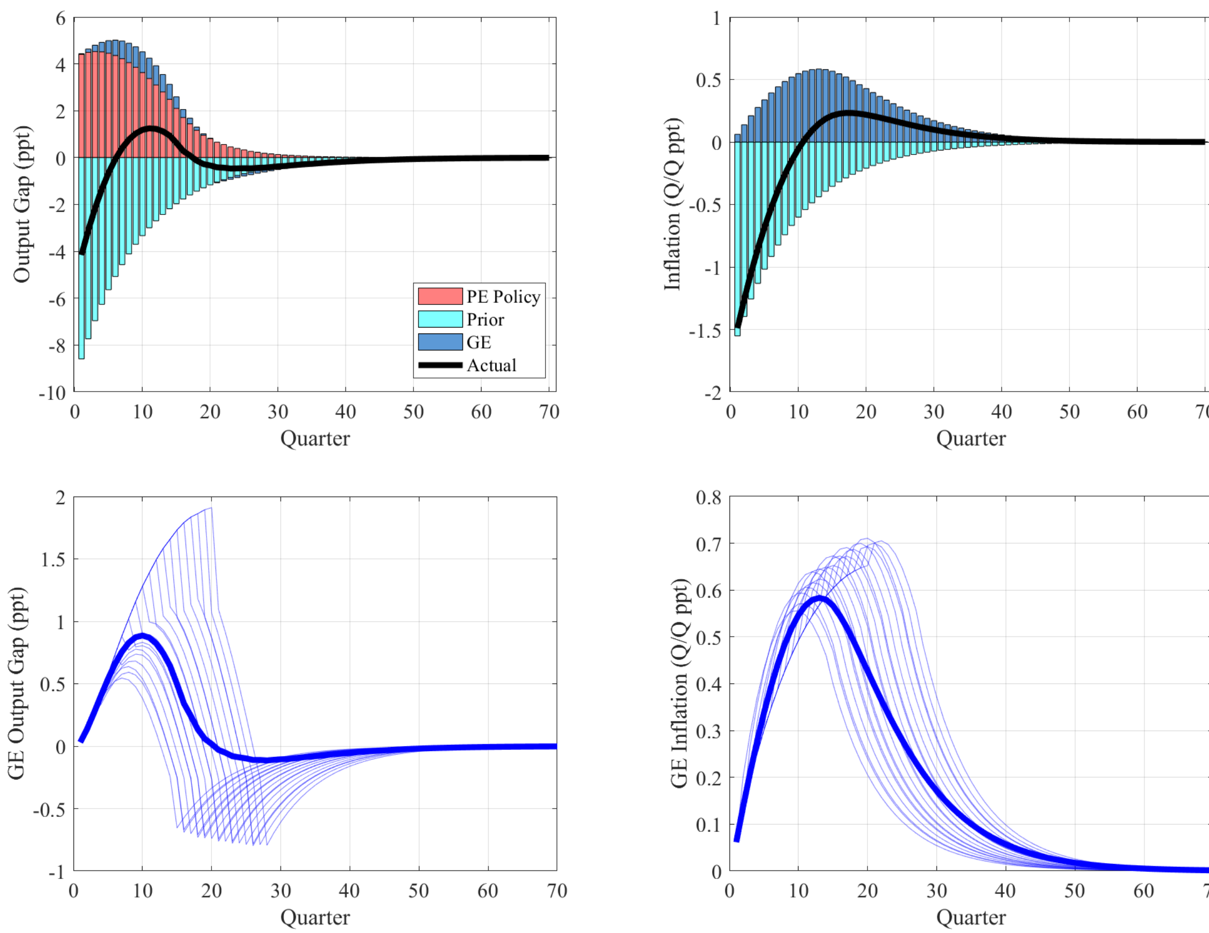

Notes: Parameter values $g = 0.15$, $\rho = 0.95$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.02$, $\lambda_x = 0.05$, $r_L = -0.003$. We plot only the first 20 realizations of uncertainty for the bottom panels for visual clarity.

Over time, even as the stimulus from forward guidance is gradually withdrawn, the general equilibrium effects tend to grow. The large initial stimulus to the output gap gradually increases inflation and inflation expectations through the Phillips curve, driving the delayed and persistent behavior of inflation under optimal policy. And while smaller on average than the partial equilibrium effects of policy, the general equilibrium effects on output are economically important, as discussed further below. The appendix demonstrates that the size of general equilibrium effects depend on the parameters that shape individual beliefs and the specification of post-crisis policy. For example, larger general equilibrium effects are associated both with higher Kalman gain coeffiecients, because a given forecast error leads to a larger markup in beliefs, and also more accommodative monetary policy on normalization, because individuals anticipate less aggressive interest rate rises in response to increasing long-term output and inflation expectations.[20]

The bottom panel of Figure 2 plots the general equilibrium effects conditional on each realization of uncertainty. Conditional on remaining in the low state, the general equilibrium effects on output provide non-trivial and gradually increasing support to the economy, as forward guidance promises are reduced. Similarly, the later is the return to the high state, the larger role they play in shaping the economy's post crisis outcomes. Not surprisingly, general equilibrium effects entirely shape the delay and persistence of inflation. Finally, general equilibrium effects of policy on output and inflation vanish in the long run, even conditional on remaining in the low state, with the economy converging to the no-forward-guidance equilibrium.

Why front-loaded promises? The necessity of front-loaded promises to manage back-loaded general equilibrium effects is well-illustrated by a counter example—an economy with a back-loaded forward guidance policy. We assume that the central bank promises no forward guidance for small shock durations, but progressively increases the forward guidance commitments each quarter to a maximum of 16 quarters, at which value it remains for all longer duration crises.

From Figure 3 several consequences are immediate. First, the average expected outcomes are considerably worse in the short-run, with a larger contraction in real activity, as well as the medium run, displaying a later peak when compared to the front-loaded policy. Second, this delayed recovery reflects the fact that general equilibrium effects of policy take even longer to kick in under the back-loaded policy, peaking on average after some 20 quarters for both output and inflation. Third, the bottom panels reveal these averages mask substantial variation in outcomes for specific realizations of uncertainty. For crises lasting longer than 10 quarters, the general equilibrium effects on output and inflation are substantial, leading to increases in real activity of greater than 2 percent, with rapidly growing effects for longer-

[20]We explore this in the appendix.

Figure 3: PE vs GE Effects
Forward guidance policy $k_\tau = 0, 1, 2, 3, ..., 15, 16, 16, ...$

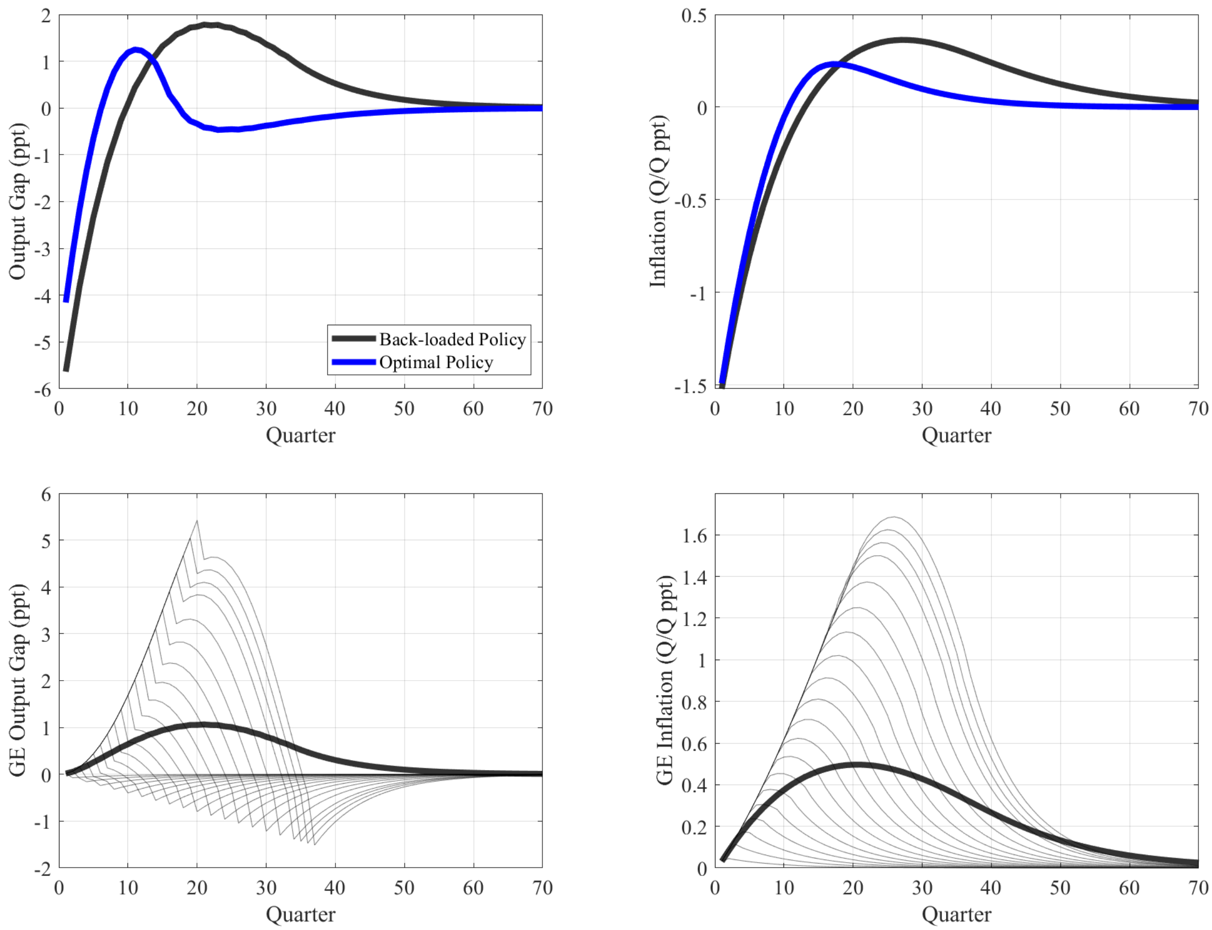

Notes: Parameter values $g = 0.15$, $\rho = 0.95$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.02$, $\lambda_x = 0.05$, $r_L = -0.003$. We plot only the first 20 realizations of uncertainty for the bottom panels for visual clarity.

duration crises.

The final two panels highlight the difficulty for stabilization policy. While short-duration shocks permit output and inflation to be stabilized quickly on return to the high state (because general equilibrium effects and forward guidance commitments are negligible), long-duration crises display considerably worse stabilization outcomes. This is evidenced by the fact that conditional on remaining in the low state, the general equilibrium effects on output and inflation continue to rise, consistent with the simple example. But the timing is wrong: these rising general equilibrium effects are further fueled by the delivery of long periods of zero interest rate policy, well after the expected end of the crisis. Long and variables lags make this policy strategy ineffective.

Robustness. The appendix demonstrates these conclusions are robust to alternative choices for the agents' learning process. Through the paper we maintain the assumption that prior beliefs $\omega_{0,L}$ are consistent with the no-forward guidance rational expectations equilibrium. This choice is inconsequential for the results in the paper. Having prior beliefs that are more optimistic or pessimistic about the aggregate effects of the shocks relative to rational expectations does not alter the key feature of the optimal monetary policy. The perceived persistent of unobserved variable ω_t affects the size and timing of the policy-induced general equilibrium effects. Higher values of the ρ parameter delays the general equilibrium effects to later periods, inducing a larger downturn after lift-off. Perhaps not surprisingly, the learning gain $\bar{g}$ affects the overall size of general equilibrium effects by making expectations more responsive to forecast errors. While the economy's response to the shock is affected by specific parameters choices, the main conclusions remain unchanged. The optimal monetary policy response is front-loaded and produces an overshooting of output and inflation in response to gradual general equilibrium effects.[21]

6.4 Comparison to Rational Expectations Equilibrium

To complete our understanding of the optimal forward guidance policy under learning, we compare the results to the optimal policy results under rational expectations in Eggertsson and Woodford (2003). Figure 4 plots the average expected dynamics under each belief assumption. It is immediate that under learning, optimal forward guidance policy must accept larger fluctuations in output and inflation, relative to a rational expectations analysis. Monetary policy is less effective. Indeed, under rational expectations, policy can largely stabilize inflation and output in the face of the same negative demand shock. The fourth panel highlights the fundamentally different shapes for forward guidance profiles under each belief assumption: large front-loaded promises are required under learning, while modest back-loaded promises are required under rational expectations.

These different commitments reflect the general equilibrium properties of policy under each belief assumption. The final two panels reveal that rational expectations deliver large immediate general equilibrium effects from a commitment to reflation in the future. For both variables, on impact, the general equilibrium effects under rational expectations are orders of magnitude larger than those under learning. Crucially, general equilibrium effects are *front-loaded* under rational expectations, as agents anticipate the full impact of stimulus on the economy in each state of the world. This very different policy transmission mechanism engenders radically different policy responses. Under rational expectations, the *promise* to provide additional stimulus if the shock persists is sufficient to trigger higher inflation expectations, lowering real interest rates, delivering economic revival. Under learning, a

[21] The appendix provides a range of examples supporting these claims.

Figure 4: Optimal Policy Relative RE Optimal Policy

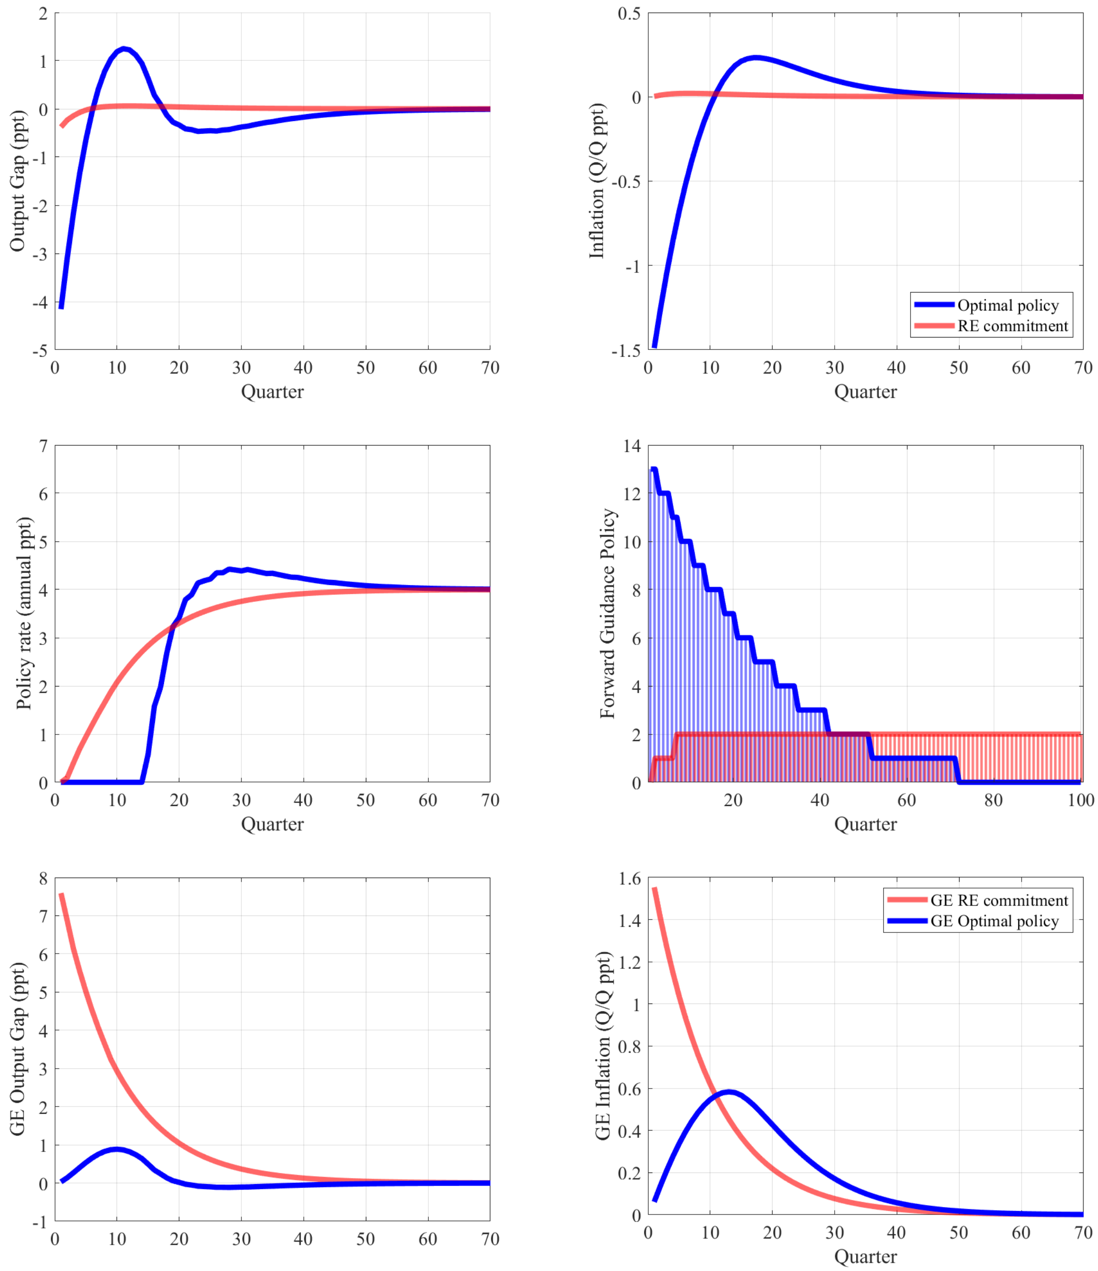

Notes: Optimal forward guidance policy (blue) compared to discretion (black) and optimal policy (red) under RE. Parameter values $g = 0.15$, $\rho = 0.95$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.02$, $\lambda_x = \kappa/7.87$, $r_L = -0.003$.

substantial upfront stimulus is needed at the beginning of the crises to ensure economic resilience to long-duration shocks.

It is worth remembering that under both belief assumptions agents fully understand state-contingent forward guidance policy announcements and the central bank is fully credible. As we argue in the Introduction, this is consistent with empirical evidence from survey interest rate forecasts. The critical difference concerns individuals' understanding of the effects of policy on the economy. Because of learning, the impact of policy on agents' output and inflation expectations is drastically smaller than rational expectations predictions. This is, again, in line with empirical evidence on the revisions of output and inflation expectations in response to forward guidance announcements.

If individuals must learn about the general equilibrium effects of monetary policy, then the central bank faces a more difficult stabilization problem: learning engenders delayed but persistent general equilibrium effects from current policy announcements. The central bank optimally manages these long and variable lags by front-loading forward guidance stimulus.[22]

7 The Insurance Principle

The optimal policy response to the crisis under learning features an insurance principle. The risk to be insured is a demand shock with uncertain duration. At the level of the macroeconomy, the central bank self insures by making large state-contingent promises for short-duration shocks to support output and inflation expectations in the event of a long-duration shock. Maintaining the power of monetary policy is the payoff. 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 output and inflation require 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.

To evaluate the costs of such commitment *ex post* we consider two exercises. First, we calculate the welfare loss from not implementing the optimal forward guidance policy, and doing nothing. Second, 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.[23] The path of the policy rate is then optimally determined, consistent with regime 3 in the earlier commitment problem.

[22]The second key different between our analysis and that of Eggertsson and Woodford (2003) concerns interest rate policy on nomalization. They assume the central bank commits to a price level targeting rule, which generates desirable history dependence. Our analysis considers a different class of rule, but one that nonetheless generates history dependence by responding to agent expectations, which are functions of the history of observed inflation and output. Such rules can replicate the overshooting property of optimal commitment policies. See Eusepi, Giannoni, and Preston (2018) and Eusepi and Preston (2018) for discussion.

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

Figure 5: Optimal Policy vs Optimal Policy Renege vs Discretion

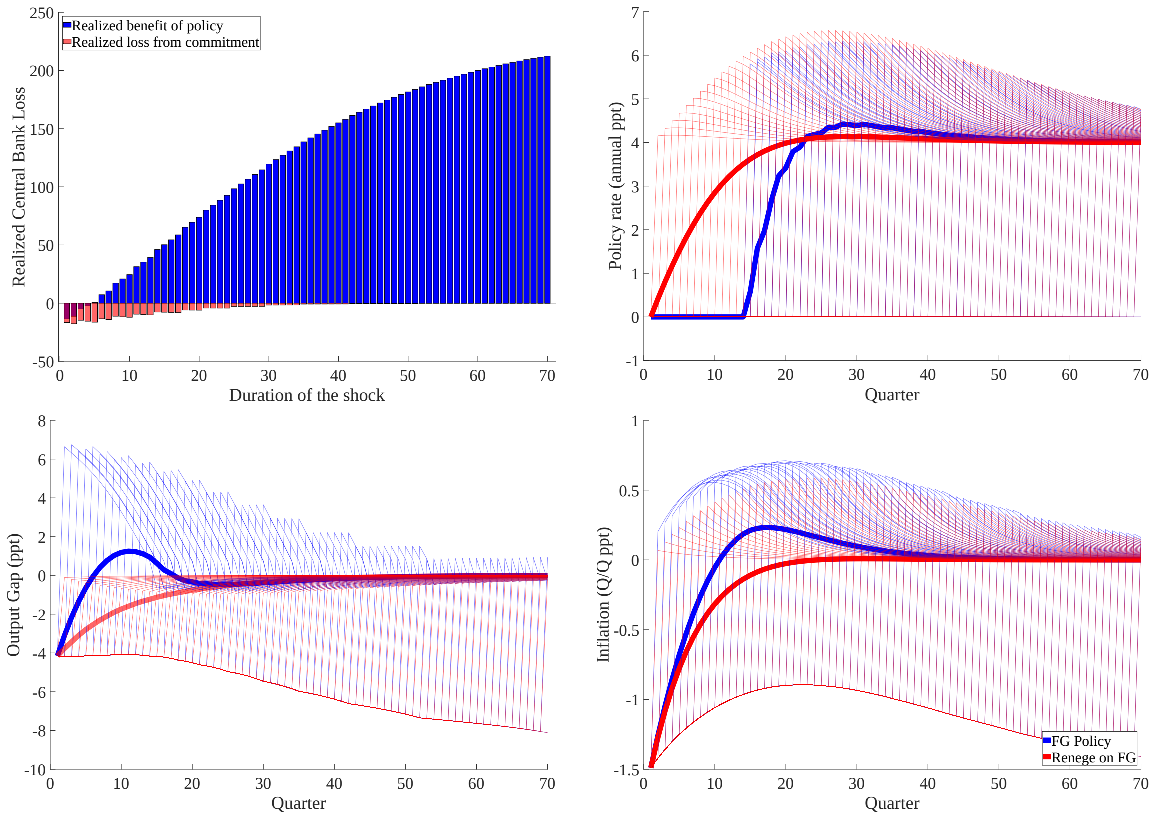

Notes: Parameter values $g = 0.15$, $\rho = 0.95$, $\beta = 0.99$, $\sigma = 0.5$, $\delta = 0.1$, $\kappa = 0.02$, $r_L = -0.003$. Bottom right figure shows a decomposition of the realized net benefit of optimal policy. The difference in central bank policy of reneging on FG policy to optimal FG policy is shown in blue.

Figure 5 displays the results. The blue bars in the first panel 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. These outcomes arise because absent forward guidance, output and inflation are stabilized in the post-crisis high state, after experiencing a significant contraction. Forward guidance ameliorates the downturn, but complicates post-crisis policy. The value of forward guidance policy reflects these competing tensions on stabilizing the economy during and after the crisis. Of course, the longer the crisis persists, the larger is the cumulative lost output from not implementing forward guidance.

That forward guidance policy can reduce welfare for highly favorable realizations of uncertainty underscores a fundamental trade-off confronting monetary policy, and the tensions between ex ante and ex post perspectives on policy. To support the economy in case of long-duration shocks 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. Said differently, if the central bank knew for sure that the crisis would be short-lived, the optimal policy would be to do nothing (at least in a choice between these two policy alternatives). But the presence of uncertainty about the duration of the shock, with the possibility of a long-duration crisis, makes immediate action desirable. In effect, the central bank takes out insurance at the level of the aggregate economy: it pays a premium in the form of worse stabilization outcomes for short-dated shocks, to secure good stabilization outcomes for long-dated shocks.

Turning to the second thought experiment, the red lines in each panel show equilibrium outcomes for a central bank that reneges on the announced forward guidance. The blue lines reproduce baseline results for the optimal forward guidance with commitment. As before the solid lines give the expected dynamics conditional on each policy, while the thin lines denote specific realizations of uncertainty.

The solid lines in the final three panels show that early abandonment of zero interest rate policy has significant implications for dynamics. On average the renege policy generates less over-shooting of inflation and never generates a boom in output. Interest rates rise by less, peaking at around 4.25 per cent, below the baseline 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, when the central bank fails to fulfill its commitment, the output gap at the time of normalization progressively falls. 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.

Returning to the first panel, the red 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. The temptation to renege is largest for short-duration shocks. 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. 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.

8 Some Practical Policy Implications

The challenges of the global pandemic birthed new debates and revisited old intellectual battle grounds about macroeconomic stabilization policy. In regards to monetary policy, two central debates concerned:

i. The efficacy and design of stimulus policy, specifically discussions about the relative merits of state-contingent versus calendar-based forward guidance policy, as well as more fundamental questions about the usefulness of such policies; and

ii. The appropriate time to normalize monetary policy when recovery appeared assured. Robust inflation was taken by critics of monetary policy to signal that central banks were either 'behind the curve' or had simply misjudged the stimulus required by economic developments, and in some cases both. And central banks themselves agonized about raising interest rates, given prior commitments to a longer period of zero interest rate policy. Would reneging on these promises lead to inferior macroeconomic control in the next crisis because of reputational concerns?

While the analysis that we present here is not intended to be a formal account of policy during the pandemic—if only because of the absence of either fiscal policy or economy-wide supply disruptions—it nonetheless displays a number of properties that give valuable perspective on the above debates, not to mention macroeconomic reality.

In early 2020 policymakers rightly thought they faced a challenge comparable to the Great Depression. None could have foreseen relatively low case fatality rates, at least compared to the grim daily news from New York, Italy and elsewhere, and the impressively quick development of vaccines that permitted resumption of normal economic life, with the reopening of sectors of the economy that were closed to protect public health. The shock was clearly of uncertain magnitude and duration.

But this "favorable realization of uncertainty" does not invalidate the substantial stimulus put in place to support the economy at the outset of the pandemic. To argue that inflation signals excessive stimulus requires an unreasonable degree of certitude that the crisis would moderate when it did. Such ex post rationalizations of policy outcomes do injustice to those making ex ante decisions in environments of great uncertainty. That is not to say we should not debate and evaluate the merits of various policy actions, but high inflation itself is no criterion to form judgments one way or the other—indeed, as our analysis makes clear, an absence of inflation during economic recovery would indicate either ineffective policy design or implementation. Similarly, that central bankers themselves desired to raise interest rates earlier than promised is not necessarily evidence of bad policy. Rather it might simply reflect the high temptation to renege on commitments in favorable circumstances.

The analysis also makes clear that monetary policy is less powerful than a rational expectations analysis might suggest, at least in situations in which individuals are unfamiliar. And the distinction between state-contingent and calendar-based forward guidance is perhaps less important than many make it out to be. But whatever one's views about the challenges of communicating central bank intentions for zero interest rate policy, building a narrative about the optimal management of large negative demand shocks is complicated by long and variables lags that stem from delayed but persistent general equilibrium effects of policy. At the onset of the crisis, the central bank has fairly tight control of inflation and output. With general equilibrium effects of policy negligible, the partial equilibrium effects under direct control of policy largely determine the degree to which the central bank mitigates the demand shock. Over time, general equilibrium effects of policy become more potent, creating non-trivial trade-offs for policy design. Observing high inflation, critics might charge that late bad management of the crisis sacrificed early good management, alleging that central banks mismanaged interest rate normalization. But this mistakes the cause of high inflation, which is the insurance premium against much worse realizations of uncertainty.

By way of conclusion, with an intellectual framework in hand and these reflections in mind, we return to where we started with Jerome Powell (2021):

> 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.

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, July). Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics 11*(3), 1–29.

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

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

Angeletos, G.-M. and C. Lian (2023). Dampening general equilibrium: incomplete information and bounded rationality. In R. Bachmann, G. Topa, and W. van der Klaauw (Eds.), *Handbook of Economic Expectations*, pp. 613–645. Academic Press.

Bilbiie, F. O. (2019, October). Optimal Forward Guidance. *American Economic Journal: Macroeconomics 11*(4), 310–345.

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.

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.

Christiano, L., M. Eichenbaum, and B. K. Johansen (2024). Slow learning. Working paper.

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

Crump, R. K., S. Eusepi, E. Moench, and B. Preston (2021, November). 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.

Eggertsson, G., S. Egiev, A. Lin, J. Platzer, and L. Riva (2021, July). A Toolkit for Solving Models with a Lower Bound on Interest Rates of Stochastic Duration. *Review of Economic Dynamics 41*, 121–173.

Eggertsson, G. and M. Woodford (2003). The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 139–211.

Eusepi, S. (2010, 03). Central Bank Communication and the Liquidity Trap. *Journal of Money, Credit and Banking 42*(2-3), 373–397.

Eusepi, S., M. P. Giannoni, and B. Preston (2018). Some implications of learning for price stability. *European Economic Review 106*, 1–20.

Eusepi, S. and B. Preston (2018). The science of monetary policy: An imperfect knowledge perspective. *Journal of Economic Literature 56*(1), 3–61.

Evans, G. W., C. Gibbs, and B. McGough (2021). A Unified Model of Learning to Forecast. forthcoming *American Economic Journal: Macroeconomics*, University of Sydney.

Evans, G. W., E. Guse, and S. Honkapohja (2008, November). 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 D. Cobham, O. Eitrheim, S. Gerlach, and J. F. Qvigstad (Eds.), *Twenty Years of Inflation Targeting: Lessons Learned and Future Prospects*, 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, November). Monetary Policy, Bounded Rationality, and Incomplete Markets. *American Economic Review 109*(11), 3887–3928.

Friedman, M. (1960). *A Program for Monetary Stability*. New York, Fordham University Press.

Gabaix, X. (2020, August). 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, January). Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis. *American Economic Review 109*(1), 86–120.

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

Marcet, A. and T. J. Sargent (1989). Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy*, 1306–1322.

Preston, B. (2005). Learning about monetary policy rules when long-horizon expectations matter. *International Journal of Central Banking 1*(2), 81–126.

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*, 1–12.

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

Woodford, M. (2013, 05). Macroeconomic Analysis Without the Rational Expectations Hypothesis. *Annual Review of Economics 5*(1), 303–346.

Woodford, M. and Y. Xie (2019, May). Policy options at the zero lower bound when foresight is limited. *AEA Papers and Proceedings 109*, 433–37.