Risk Management for Monetary Policy
Near the Zero Lower Bound[*]

Charles Evans   Jonas Fisher   François Gourio   Spencer Krane

March 13, 2015

Abstract

As projections have inflation heading back toward target and the labor market continuing to improve, the Federal Reserve has begun to contemplate an increase in the federal funds rate. There is however substantial uncertainty around these projections. How should this uncertainty affect monetary policy? In standard models uncertainty has no effect. In this paper, we demonstrate that the zero lower bound on nominal interest rates implies that the central bank should adopt a looser policy when there is uncertainty. In the current context this result implies that a delayed liftoff is optimal. We first demonstrate this result theoretically in two canonical macroeconomic models. On the one hand, raising rates early might lead to low output and inflation if the economic fundamentals turn out weaker than expected. On the other hand, raising rates later might lead to high inflation if economic fundamentals are stronger than expected. Near the zero lower bound, monetary policy tools are strongly asymmetric and can deal with the second scenario much more easily than with the first. We next provide a quantitative evaluation of this policy using numerical simulations calibrated to the current environment. Finally, we present narratives from Federal Reserve communications that suggest risk management is a longstanding practise, and econometric evidence that the Federal Reserve historically has responded to uncertainty, as measured by a variety of indicators.

JEL Classification Numbers: E3, E4, E5, E6
Keywords: monetary policy, risk management, zero lower bound

---

1 Introduction

To what extent should uncertainty affect monetary policy? This classic question is relevant today as the Fed considers when to start increasing the federal funds rate (FFR). In the December 2014 Summary of Economic Projections (SEP), most Federal Open Market Committee (FOMC) participants forecast that the unemployment rate would return to its long-run neutral level by late 2015 and that inflation would gradually rise back to its 2 percent target. This forecast could go wrong in two ways. One is that the FOMC may be overestimating the underlying strength in the real economy or in the factors that will return inflation to target. Guarding against these risks calls for cautious removal of accommodation. The second is that we could be poised for a much stronger rise in inflation than currently projected. This risk calls for more aggressive rate hikes. How should policy manage these divergent risks?

In our view, the biggest risk we face today is prematurely engineering restrictive monetary conditions. If the FOMC misjudges the impediments to growth and inflation and reduces monetary accommodation too soon, it could find itself in the very uncomfortable position of falling back into the ZLB environment. The implications of the ZLB for the attainment of the FOMC's policy goals are severe. It is true the FOMC has access to unconventional policy tools at the ZLB, but these appear to be imperfect substitutes for the traditional FFR instrument. Furthermore, there is no guarantee that they will be as successful as they have in the past if monetary policy tightened and then the economy were to soon return to the ZLB, in part because the credibility that supported the alternative tools' prior efficacy could be substantially diminished by an unduly hasty exit from the ZLB.

In contrast, it is reasonable to imagine that the costs of inflation running moderately above target for a while are much smaller than the costs of falling back into the ZLB. This is not the least because it is likely that inflation could be brought back into check with modest increases in interest rates. These measured rate increases likely would be manageable for the real economy, particularly if industry and labor markets had already overcome the headwinds that have kept productive resources from being efficiently and fully employed. In addition,

inflation in the U.S. has averaged well under that 2 percent mark for the past six and a half years. With a symmetric inflation target, one could imagine moderately-above-target inflation for a limited period of time as simply the flip side of the recent inflation experience – and hardly an event that would impose great costs on the economy.

To summarize, raising rates too early increases the likelihood of adverse shocks driving the economy back to the ZLB, while delaying lift-off too long against a more robust economy could lead to an unwelcome bout of high inflation. Since the tools available to counter the first scenario may be less effective than the traditional tool of raising rates to counter the second scenario, the costs of premature lift-off exceed those of delay. It therefore seems prudent to refrain from raising rates until we are highly certain that the economy has achieved a sustained period of strong growth and that inflation is on a clear trajectory to return to target.[1]

In this paper we establish theoretically that in the current setting uncertainty about monetary policy being constrained by the ZLB in the future implies an optimal policy of delayed lift-off–the risk management framework just described. We formally define risk management as the principle that policy should be formulated taking into account the dispersion of shocks around their means. Our main theoretical contribution is to provide a simple demonstration that within the canonical framework used to study optimal monetary policy under discretion, the ZLB implies a new role for such risk management through two distinct economic channels.

The first channel - which we call the *expectations channel*– arises because the possibility of a binding ZLB tomorrow leads to lower expected inflation and output today, and hence requires some counteracting policy easing today. The second channel – which we call the *buffer stock* channel – arises because it can be useful to build up output or inflation today in order to reduce the likelihood and severity of hitting the ZLB tomorrow. We show that optimal policy when one of these channels is operative dictates that lift-off from a zero interest rate should be delayed at times when a return to the ZLB remains a distinct possibility. These

---

[1]Evans (2014)'s speech at the Petersen Institute of Economics discusses these issues at greater length.

channels operate in very standard macroeconomic models, so no leap of faith is necessary to embrace them, at least at a qualitative level. We use simulations of our theoretical models calibrated to the current environment to determine the quantitative relevance of the two channels. For plausible degrees of uncertainty we find optimal policy prescribes 3 to 6 quarters of delay in liftoff from the ZLB relative to a popular policy reaction function that does not take into account this uncertainty.

While we establish a solid theoretical basis for risk management near the ZLB, it is natural to ask whether such a framework is relevant for policy away from the ZLB. The answer is yes. It is true that in a wide class of models that abstract from the ZLB, optimal policy typically involves adjusting the interest rate in response to the mean of the distribution of shocks and information on higher moments is irrelevant. However, there also is an extensive literature covering departures from this result based on nonlinear economic environments or uncertain policy parameters that justify taking a risk management approach away from the ZLB. This brings us to the question of whether or not policy-makers have actually taken this approach in the past. Is risk management old hat for the FOMC? We explore this question in two ways.

First, we analyze monetary policy communications over the period pre-ZLB period 1993-2008 and find evidence that risk management has been a long-standing operating characteristic of the FOMC, at least in words if not in deeds. In particular, we find numerous examples when uncertainty and insurance have been used to explain monetary policy settings. This analysis demonstrates that calling for a risk management approach in the current policy environment is not out of the ordinary and in fact is a well-established approach to monetary policy. Confirmation of this view is found in Greenspan (2004) who states ". . . the conduct of monetary policy in the United States has come to involve, at its core, crucial elements of risk management."

Second, we explore whether the words of the FOMC are reflected in policy actions prior to the ZLB period. For this analysis we estimate a conventional policy reaction function augmented with a variety of measures of risk and test if these additional terms are signifi-

cantly different from zero. The measures of risk we look at include ones based on financial market data, revisions to Federal Reserve Board staff forecasts, survey measures of forecasts, and several measures derived from a narrative analysis of the FOMC minutes. We find clear evidence that risk in the economic outlook has had a material impact on the interest rate choices of the FOMC above and beyond point forecasts of inflation and the output gap.

The remainder of this introduction discusses why we think unconventional policies are unlikely to be good substitutes for interest rate policies when they are constrained by the ZLB. Our theoretical analysis is predicated on the simplifying assumption that the only policy instrument available is control over short term interest rates. If the monetary policy toolkit contained alternative instruments that were perfect substitutes for changing the policy rate, then the ZLB would not present any special economic risk and our analysis would be moot. However, even though most central bankers believe unconventional policies such as large scale asset purchases (LSAPs) or forward guidance about policy rates can provide considerable accommodation at the ZLB, no one argues that these tools are on an equal footing with traditional policy instruments.[2]

One reason for this is that effects of unconventional policies on the economy naturally are much more uncertain than the transmission mechanism from traditional tools. Various studies of LSAPs, for example, have generated a wide range of estimates of their ability to put downward pressure on private borrowing rates and influence the real economy. Furthermore, the effects of both LSAPs and forward guidance about the future path of the FFR are complicated functions of private sector expectations, which makes their economic effects highly uncertain.[3]

Alternative monetary policy tools also carry potential costs that are somewhat different

---

[2]For example, while there is econometric evidence that changes in term premia influence activity and inflation, the effects appear to be less powerful than comparably sized movements in the short term policy rate, see D'Amico and King (2015), Kiley (2012) and Chen, Curida, and Ferrero (2012).

[3]Bomfin and Meyer (2010), D'Amico and King (2013) and Gagnon, Raskin, Remache, and Sack (2010) find noticeable effects of LSAPs on Treasury term permia while Chen et al. (2012) and Hamilton and Wu (2010) unearth only small effects. Krishnamurthy and Vissing-Jorgensen (2013) argue that the LSAPs have only had a substantial influence on private borrowing rates in the mortgage market. Engen, Laubach, and Reifschneider (2015) and Campbell, Evans, Fisher, and Justiniano (2012) provide analysis of the dependence of LSAPs and forward guidance on private sector expectations.

from those associated with standard policy. The four most commonly cited costs are: the large increases in reserves generated by LSAPs risk unleashing inflation; a large balance sheet may make it more difficult for the Fed to raise interest rates when the time comes; the extended period of low interest rates and Federal Reserve intervention in the long-term Treasury and mortgage markets may induce inefficient allocation of credit and financial fragility; and the large balance sheet puts the Federal Reserve at risk of incurring financial losses if rates rise too quickly and such losses could undermine its support and independence.[4] For the most part these costs appear to be very hard to quantify, and so naturally elevate the level of uncertainty associated with ZLB policies.

A consequence of all of the uncertainty over benefits and costs is that unconventional tools are likely to be used more cautiously than traditional policy instruments. For example Bernanke (2005) emphasizes that because of the uncertain costs and benefits of them ". . . the hurdle for using unconventional policies should be higher than for traditional policies." Furthermore, at least conceptually, some of the benefits of ZLB policies may be decreasing, and the costs increasing, in the size of the balance sheet or in the amount of time spent in a very low interest rate environment.[5] Accordingly, policies that had wide-spread support early on in a ZLB episode might be difficult to extend or expand with an already large balance sheet or with smaller shortfalls in policy targets.

So, while valuable, alternative policies also appear to be less-than-perfect substitutes for changes in short term policy rates. Accordingly, the ZLB presents a different set of risks to policymakers than those that they face during more conventional times and thus it is worthy of consideration on its own accord. We abstract from unconventional policy tools for the remainder of our analysis.

---

[4]These costs are mitigated, however, by additional tools the Fed has introduced to enhance control over interest rates when the time comes to exit the ZLB and by enhanced supervisory and regulatory efforts to monitor and address potential financial stability concerns. Furthermore, continued low rates of inflation and contained private-sector inflationary expectations have reduced concerns regarding an outbreak of inflation.

[5]Krishnamurthy and Vissing-Jorgensen (2013) argue successive LSAP programs have had a diminishing influence on term premia. Surveys conducted by Blue Chip and the Federal Reserve Bank of New York also indicate that market participants are less optimistic that further asset purchases would provide much stimulus if the Fed was forced to expand their use in light of unexpected economic weakness.

2 Rationales for Risk Management Near the ZLB

The canonical framework of monetary policy analysis assumes that the central bank sets the nominal interest rate to minimize a quadratic loss function of the deviation of inflation from its target and the output gap, and that the economy is described by a set of linear equations. This framework allows the optimal interest to be calculated as a function of the underlying shocks or economic fundamentals. In most applications, uncertainty is incorporated as additive shocks to these linear equations capturing factors outside the model that lead to variation in economic activity or inflation.[6]

A limitation of this approach is that, by construction, it denies that a policymaker might choose to adjust policy in the face of changes in uncertainty about economic fundamentals. However, the evidence discussed below in Section 3 suggests that in practice, policymakers are sensitive to uncertainty and respond by following what appears to be a risk management approach. Motivating why a central banker should behave in this way requires some departure from the canonical framework. The main contribution of this section is to consider a departure associated with the possibility of a binding ZLB in the future.

We show that when a policymaker might be constrained by the ZLB in the future, optimal policy today should take account of uncertainty about fundamentals. We focus on two distinct channels through which this can occur. To keep the analysis transparent we study these channels using two closely related but different models. We first use the workhorse forward-looking New Keynesian model to illustrate the *expectations channel*, in which the possibility of a binding ZLB tomorrow leads today to lower expected inflation and output gap, thus necessitating policy easing today. We then use a backward-looking "Old" Keynesian set-up to illustrate the *buffer stock channel*, in which it can be optimal to build up output or inflation today in order to reduce the likelihood and severity of being constrained by the ZLB tomorrow.[7] After describing these two channels we study some

---

[6]This framework can be derived from a micro-founded DSGE model (see for instance Woodford (2003), Chapter 6), but it has a longer history and is used even in models that are not fully micro-founded. For instance, the Federal Reserve Board staff routinely conducts optimal policy exercises in the FRB-US model, see for example English, López-Salido, and Tetlow (2013)

[7]Both of these channels operate in modern DSGE models such as Christiano, Eichenbaum, and Evans

numerical simulations to assess their quantitative effects.

2.1 The Expectations Channel

The simple New Keynesian model has well established micro-foundations based on price stickiness. Given that there are many excellent expositions of these foundations, *e.g.* Woodford (2003) or Gali (2008), we just state our notation without much explanation. The model consists of two main equations, the *Phillips curve* and the *IS curve*.

The Phillips curve is specified as

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t. \tag{1}$$

In (1) π_t and x_t are both endogenous variables and denote inflation and the output gap at date t. E_t is the date t conditional expectations operator; rational expectations is assumed. The variable u_t is a mean zero exogenous *cost-push* shock, $0 < \beta < 1$, and $\kappa > 0$. For simplicity we assume the central bank has a constant inflation target equal to zero so π_t is the deviation of inflation from that target.[8] The cost-push shock represents exogenous changes to inflation such as an independent decline in inflation expectations, dollar appreciation or changes in oil prices.

The IS curve is specified as

$$x_t = E_t x_{t+1} - \frac{1}{\sigma}\left(i_t - E_t \pi_{t+1} - \rho_t^n\right), \tag{2}$$

where $\sigma > 0$, i_t is the nominal interest rate controlled by the central bank, and ρ_t^n is the *natural rate of interest* given by

$$\rho_t^n = \rho + \sigma g_t + \sigma E_t(z_{t+1} - z_t). \tag{3}$$

In this specification of the natural rate g_t is a mean zero *demand* shock, and z_t is the exogenous log of potential output. With constant potential output and demand shock equal to zero the natural rate equals $\rho > 0$.

The natural rate is the setting of the nominal interest rate consistent with expected inflation at target and the output gap equal to zero.[9] Since z_t and g_t are exogenous, so is the natural rate. Our analysis is centered around uncertainty in the natural rate. From (3) this arises due to uncertainty about g_t and $E_t(z_{t+1} - z_t)$. We interpret the former as arising due to a variety of factors, including fiscal policy, foreign economies' growth, and financial considerations such as de-leveraging.[10] Clearly the latter source of uncertainty is over the variety of factors that can influence the expected rate of growth in potential output, for example as emphasized in the recent debate over "secular stagnation."

We adopt the canonical framework in assuming the central bank acts to minimize a quadratic loss function with the understanding that private-sector behavior is governed by (1)–(3). The loss function is

$$L = \frac{1}{2} E_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda x_t^2 \right), \tag{4}$$

where $\lambda \geq 0$. We further assume the ZLB constraint, *i.e.* $i_t \geq 0$. We abstract from unconventional policies and assume the short term interest rate is the central bank's only policy instrument. It is set by solving for optimal policy under *discretion*. In particular, each period the central bank sets the nominal interest rate given the current state with the understanding that private agents anticipate that the central bank will re-optimize in the following periods.

---

[9]Woodford (2003, p. 248) defines the natural rate as the equilibrium real rate of return in the case of fully flexible prices. As discussed by Barsky, Justiniano, and Melosi (2014), in medium-scale DSGE models with many shocks the appropriate definition of the natural rate is less clear.

[10]Uncertainty itself could give rise to g_t shocks. A large amount of recent work, following Bloom (2009), suggests that private agents react to increases in economic uncertainty, leading to a decline in economic activity. One channel is that higher uncertainty may lead to precautionary savings which depresses demand, as in emphasized by Basu and Bundick (2013), Fernández-Villaverde, Guerró-Quintana, Kuester, and Rubio-Ramírez (2012) and Born and Pfeifer (2014).

We focus on optimal policy under discretion for two reasons. First, the case of commitment with a binding ZLB already has been studied extensively. In particular it is well known from the contributions of Krugman (1998), Egertsson and Woodford (2003), Woodford (2012) and Werning (2012) that commitment can reduce the severity of the ZLB problem by creating higher expectations of inflation and the output gap. One implication of these studies is that the central bank should commit to keeping the policy rate at zero longer than would be prescribed by discretionary policy. By studying optimal policy under discretion we find a different rationale for a policy of keeping rates "lower for longer" that does not rely on the central bank having the ability to commit to a time-inconsistent policy.[11] Second, we think this approach better approximates the institutional environment in which the FOMC operates.

2.1.1 A ZLB Scenario

We study optimal policy when the central bank is faced with the following simple ZLB scenario. The central bank observes the current value of the natural rate, ρ_0^n, and the cost-push shock u_0; moreover, there is no uncertainty in the natural rate after $t = 2$, $\rho_t^n = \rho > 0$ for all $t \geq 2$, nor in the cost push shock after $t = 1$, $u_t = 0$ for all $t \geq 1$. However, there is uncertainty at $t = 1$ regarding the natural rate ρ_1^n.[12] The variable ρ_1^n is assumed to be distributed according to the probability density function $f_\rho(\cdot)$.[13]

This very simple scenario keeps the optimal policy calculation tractable while preserving the main insights. We also think it captures some key elements of uncertainty faced by the FOMC today. Since (1) and (2) do not contain endogenous state variables we do not have to take a stand on whether the ZLB is binding before $t = 0$. One possibility is that the

---

[11]Implicitly we are assuming the central bank does not have the ability to employ what Campbell et al. (2012) call "Odyssean" forward guidance. However our model is consistent with the central bank using forward guidance in the "Delphic" sense they describe because agents anticipate how the central bank reacts to evolving economic conditions.

[12]It is easy to verify that if the uncertainty about the natural rate is only at $t = 0$ the optimal policy would be to set the interest rate to the expected value of the natural rate and and the amount of uncertainty would have no affect. This is why our scenario has more than two periods.

[13]There is ample evidence of considerable uncertainty regarding the natural rate. See for example Barsky et al. (2014), Hamilton, Harris, Hatzius, and West (2015) and Laubach and Williams (2003).

natural rate ρ_t^n is negative for $t < 0$ and that the policy rate is set at zero, $i_t = 0$ for $t < 0$, but by $t = 0$ the economy is close to being unconstrained by the ZLB. However, there is uncertainty as to when this will happen. The natural rate might be low enough at $t = 1$ such that the ZLB still binds, or the economy may recover enough so that the natural rate is positive. This allows us to consider the optimal timing of lift-off.

2.1.2 Analysis

To find the optimal policy, we solve the model backwards from $t = 2$ and focus on the policy choice at $t = 0$. First, for $t \geq 2$, it is possible to perfectly stabilize the economy by setting the nominal interest rate equal to the (now positive) natural rate, $i_t = \rho_t^n = \rho$. This leads to $\pi_t = x_t = 0$ for $t \geq 2$.[14] The optimal policy at $t = 1$ will depend on the realized value of the natural rate ρ_1^n. If $\rho_1^n \geq 0$, then it is again possible (and optimal) to perfectly stabilize by setting $i_1 = \rho_1^n$, leading to $x_1 = \pi_1 = 0$. However if $\rho_1^n < 0$, the ZLB binds and consequently $x_1 = \rho_1^n/\sigma < 0$ and $\pi_1 = \kappa\rho_1^n/\sigma < 0$. The expected output gap at $t = 1$ is hence $E_0 x_1 = \int_{-\infty}^{0} \rho f_\rho(\rho)d\rho/\sigma \leq 0$ and expected inflation is $E_0\pi_1 = \kappa E_0 x_1 < 0$.

Because agents are forward-looking, this low expected output gap and inflation feed backward to $t = 0$. A low output gap tomorrow depresses output today by a wealth effect via the IS curve and depresses inflation today through the Phillips curve. Low inflation tomorrow depresses inflation since price setting is forward looking in the Phillips curve and depresses output today by raising the real interest rate via the IS curve.[15] The optimal policy at $t = 0$ must take into account these effects. This implies that optimal policy will be looser than if there was no chance that the ZLB binds tomorrow.

Substituting for π_0 and i_0 using (1) and (2), and taking into account the ZLB constraint,

---

[14]This simple interest rate rule implements the equilibrium $\pi_t = x_t = 0$, but is also consistent with other equilibria. However there are standard ways to rule out these other equilibria. See for instance Gali (2008, pp. 76–77) for a discussion. From now on, we will not mention this issue.

[15]Johannsen (2014) and Nakata (2013b) describe this mechanism in economies with an interest rate feedback rule. Note that it is distinct from the precautionary savings motives at the ZLB discussed in footnote 10.

optimal policy at $t = 0$ solves the following problem:

$$\min_{x_0} \frac{1}{2} \left((\kappa x_0 + \beta E_0 \pi_1 + u_0)^2 + \lambda x_0^2 \right) \text{ s.t. } x_0 \leq E_0 x_1 + \frac{1}{\sigma} (\rho_0^n + E_0 \pi_1).$$

Two cases arise, depending on whether the ZLB binds at $t = 0$ or not. Define the threshold value

$$\rho_0^* = -\sigma \frac{\kappa}{\lambda + \kappa^2} u_0 - \left(1 + \frac{\kappa}{\sigma} + \beta \frac{\kappa^2}{\lambda + \kappa^2} \right) \int_{-\infty}^{0} \rho f_\rho(\rho) d\rho.$$

If $\rho_0^n > \rho_0^*$, then the optimal policy is to follow the standard monetary policy response to an inflation shock to the Phillips curve, $\beta E_0 \pi_1 + u_0$ leading to:

$$x_0 = -\frac{\kappa}{\lambda + \kappa^2} (\beta E_0 \pi_1 + u_0); \quad \pi_0 = \frac{\lambda}{\lambda + \kappa^2} (\beta E_0 \pi_1 + u_0).$$

The corresponding interest rate is

$$\begin{aligned} i_0 &= \rho_0^n + E_0 \pi_1 + \sigma (E_0 x_1 - x_0), \\ &= \rho_0^n + \sigma \frac{\kappa}{\lambda + \kappa^2} u_0 + \left(1 + \frac{\kappa}{\sigma} + \beta \frac{\kappa^2}{\lambda + \kappa^2} \right) \int_{-\infty}^{0} \rho f_\rho(\rho) d\rho. \end{aligned} \tag{5}$$

As long as $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho < 0$, (5) implies that the optimal interest rate is lower than if there was no chance of a binding ZLB tomorrow, $i.e.$ the case of $f_\rho(\rho) = 0$ for $\rho \leq 0$. The interest rate is lower today to offset the deflationary and recessionary effects of the possibility of a binding ZLB tomorrow. If $\rho_0^n < \rho_0^*$, then the ZLB binds today and optimal policy is $i_0 = 0$.[16] Notice from (5) that for some parameters, the ZLB will bind at $t = 0$ even though it would not bind if agents were certain that the economy would perform well afterwards. Specifically, if agents were certain that the ZLB would not bind at $t = 1$, $E_0 x_1 = E_0 \pi_1 = 0$ and $i_0 = 0$ if $\rho_0^n \leq -\sigma \kappa u_0 / (\lambda + \kappa^2)$. So the possibility of the ZLB binding tomorrow increases the chances

---

[16]Since $E_0 x_1$ is a sufficient statistic for $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho$ in (5), the optimal policy has the flavor of a traditional forward-looking policy reaction function. However $E_0 x_1$ is not independent of a mean preserving spread in the distribution of ρ_1^n so optimal policy here departs from the certainty equivalence principle which says that the extent of uncertainty *in the underlying fundamentals* does not affect the optimal interest rate. Recent statements of the certainty equivalence principle in models with forward-looking variables can be found in Svensson and Woodford (2002, 2003).

of being constrained by the ZLB today.

Turning specifically to the issue of uncertainty, we obtain the following unambiguous result:

Proposition 1 *Higher uncertainty,* i.e. *a mean-preserving spread, in the distribution of the natural rate tomorrow* ρ_1^n *leads to a looser policy today.*

To see this, rewrite the key quantity $\int_{-\infty}^0 \rho f_\rho(\rho)d\rho = E\min(\rho, 0)$. Since the min function is concave, higher uncertainty through a mean-preserving spread about ρ_1^n leads to lower, *i.e.* more negative, $E_0 x_1$ and $E_0 \pi_1$, and hence lower i_0.[17]

Another interesting feature of the solution is that the distribution of the positive values of ρ_1^n is irrelevant for policy. That is, policy is set only with respect to the states of world in which the ZLB might bind tomorrow. The logic is that if a very high value of ρ_1^n is realized, monetary policy can adjust to it and prevent a bout of inflation. This is a consequence of the standard principle that, outside the ZLB, demand shocks can and should be perfectly offset by monetary policy.

2.1.3 Discussion

Proposition 1 has several predecessors; perhaps the closest are Adam and Billi (2007), Nakata (2013a,b) and Nakov (2008) who demonstrate numerically how, in a stochastic environment, the ZLB leads the central bank to adopt a looser policy. Our contribution is to provide a simple analytical example.[18] This result has been correctly interpreted to mean that if negative shocks to the natural rate lead the economy to be close to the ZLB, the optimal response is to reduce the interest rate aggressively to reduce the likelihood that the ZLB becomes binding. The same logic applies to liftoff. Following an episode where the ZLB has been a binding constraint, the central bank should not raise rates as if the ZLB constraint

---

[17]See Mas-Colell, Whinston, and Green (1995, Proposition 6.D.2, pp. 199) for the relevant result regarding the effect of a mean preserving spread on the expected value of concave functions of a random variable.

[18]See also Nakata and Schmidt (2014) for a related analytical result in a model with two-state Markov shocks.

were gone forever. Even though the best forecast may be that the economy will recover and exit the ZLB – *i.e.* in the context of the model, that $E_0(\rho_1^n) > 0$ – it can be optimal to have zero interest rates today. Note that policy is looser when the probability of being constrained by the ZLB in the future is high or the potential severity of the ZLB problem is large, *i.e.* $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho$ is a large negative number; the economy is less sensitive to interest rates (high σ), and the Phillips curve is steep (high κ).

While we have deliberately focused on a very simple example, our results hold under much more general conditions. For instance, the same results still hold if $\{\rho_t^n\}_{t \geq 2}$ follows an arbitrary stochastic process as long as it is positive. In the appendix we consider the case of optimal policy with uncertainty about cost-push inflation. We show that in this case as well optimal policy is looser if there is a chance of a binding ZLB in the future due to a low cost-push shock. Another implication of this case is that the risk that inflation picks up due to a high cost-push shock does not affect policy today. If such a shock were to occur tomorrow, it will lead to some inflation; however, there is nothing that policy *today* can do about it. Finally, while the model chosen is highly stylized, the core insights would likely continue to hold in a medium-scale model with a variety of shocks and frictions.

One reading of these results is that lift-off from currently near zero interest rates should be delayed, but the FOMC should be prepared to raise rates quickly if the economy actually picks up strongly. Also note that the connection of current inflation to expected inflation in (1) suggests the recent decline in inflation and in measures of inflation compensation might reflect rising expectations of the ZLB binding in the future.[19] From this perspective, the relatively low inflation measures we see today justify continuing delay in lifting off from the ZLB.

There are two obvious limitations to these results. First, it requires that the central bank is able to offset shocks to the natural rate outside the ZLB, and that there is no cost to doing so. However, while this feature greatly simplifies the analysis, we do not think it is crucial for our results. Second, we have assumed that there is no cost to raising rates quickly if needed.

---

[19]In the January 2015 Federal Reserve Bank of New York survey of Primary Dealers, respondents put the odds of returning to the ZLB within two years following liftoff at 20%.

That is, our welfare criterion does not give any value to interest rate smoothing. The policy recommendation to reduce the interest rate when there is more uncertainty naturally implies that the rate will rise on average faster over time once the economy recovers.

2.2   The Buffer Stock Channel

The buffer stock channel does not rely on forward-looking behavior, but rather on the view that the economy has some inherent *momentum*, *e.g.* due to adaptive inflation expectations, inflation indexation, habit persistence, adjustment costs and hysteresis. Suppose that output or inflation have a tendency to persist. If there is a risk that the ZLB binds tomorrow, building up output and inflation today creates some buffer against hitting the ZLB tomorrow. This intuition does not guarantee that it is optimal to increase output or inflation today. In particular, the benefit of higher inflation or output today in the event that a ZLB event arises tomorrow must of course be weighed against the costs of excess output and inflation today, and tomorrow's cost to bring down the output gap or inflation if the economy turns out not to hit the ZLB constraint. So it is important to verify that our intuition holds up in a model.

To isolate the buffer stock channel from the expectations channel we focus on a purely backward-looking "Old" Keynesian model. Purely backward-looking models do not have micro-foundations like the New Keynesian model does, but backward-looking elements appear to be important empirically.[20] Backward-looking models have been studied extensively in the literature, including by Laubach and Williams (2003), Orphanides and Williams (2002), Reifschneider and Williams (2000) and Rudebusch and Svensson (1999).

The model we study replaces (1) and (2) from the forward-looking model with

$$\pi_t \;\; = \;\; \xi \pi_{t-1} + \kappa x_t + u_t; \tag{6}$$

---

[20]Indeed empirical studies based on medium-scale DSGE models, such as those considered by Christiano et al. (2005) and Smets and Wouters (2007), find backward-looking elements are essential to account for the empirical dynamics. Backward-looking terms are important in single-equation estimation as well. See for example Fuhrer (2000), Gali and Gertler (1999) and Eichenbaum and Fisher (2007).

$$x_t = \delta x_{t-1} - \frac{1}{\sigma}\left(i_t - \rho_t^n - \pi_{t-1}\right), \tag{7}$$

where $0 < \xi < 1$ and $0 < \delta < 1$. This model is essentially the same as the one studied by Reifschneider and Williams (2000). Unlike in the New Keynesian model it is difficult to map ρ_t^n directly to underlying fundamental shocks as we do in equation (3). For simplicity we continue to refer to this exogenous variable as the natural rate and use (3) as a guide to interpreting it, but it is perhaps better to think of it as simply a "demand" shock or "IS" shock. We study optimal policy under discretion in the same way as before, and in particular under the ZLB scenario described in Section 2.1.1. We again focus on the case where $u_t = 0$ for $t \geq 1$.

2.2.1 Analysis

As before, we solve the model backwards from $t = 2$ to determine optimal policy at $t = 0$ and how this is affected by uncertainty in the natural rate at $t = 1$. After $t = 1$ the economy does not experience any more shocks so $\rho_t^n = \rho$ for $t \geq 2$, but it inherits initial lagged inflation and output terms π_1 and x_1, which may be positive or negative. The output gap term can be easily adjusted by changing the interest rate i_t, provided the central bank is not constrained by the ZLB at $t = 2$, *i.e.* if $\rho_2^n = \rho$ is large enough, an assumption we will maintain.[21] Given the quadratic loss, it is optimal to smooth this adjustment over time, so the economy will converge back to its steady-state slowly. The details of this adjustment process after $t = 2$ are not very important for our analysis. What is important is that the overall loss of starting from $t = 2$ with a lagged inflation π_1 and output gap x_1 turns out to be a quadratic function of π_1 only; we can write it as $W\pi_1^2/2$, where W is a constant that depends on λ, κ, ξ and β and is calculated in the appendix.

Turn now to optimal policy at $t = 1$. Take the realization of ρ_1^n and last period's output gap x_0 and inflation π_0 as given. Substituting for π_1 and i_1 using (6) and (7), and taking

---

[21]Relaxing it would only strengthen our results.

into account the ZLB constraint, optimal policy at $t = 1$ solves the following problem:

$$V(x_0, \pi_0, \rho_1^n) = \min_{x_1} \frac{1}{2} \left((\xi\pi_0 + \kappa x_1)^2 + \lambda x_1^2 \right) + \beta \frac{W}{2} \pi_1^2 \text{ s.t. } x_1 \leq \delta x_0 + \frac{\pi_0 + \rho_1^n}{\sigma}.$$

where the policymaker now anticipates the cost of having inflation π_1 tomorrow, and her choices are affected by yesterday's values x_0 and π_0.

Depending on the value of ρ_1^n, two cases can arise. Define the threshold value:

$$\rho_1^*(x_0, \pi_0) = -\left(\frac{(1 + \beta W)\kappa\xi}{(1 + \beta W)\kappa^2 + \lambda}\sigma + 1 \right) \pi_0 - \sigma\delta x_0. \tag{8}$$

For $\rho_1^n \geq \rho_1^*(x_0, \pi_0)$ the ZLB is not binding, otherwise it is. Hence the probability of hitting the ZLB is

$$P(x_0, \pi_0) = \int_{-\infty}^{\rho_1^*(x_0, \pi_0)} f_\rho(\rho)d\rho.$$

In contrast to the forward-looking case where the probability of being constrained by the ZLB constraint is exogenous, it is now endogenous at $t = 1$ and can be influenced by policy at $t = 0$. As indicated by (8), a higher output gap or inflation at $t = 0$ will reduce the likelihood of hitting the ZLB at $t = 1$.

If $\rho_1^n \geq \rho_1^*(x_0, \pi_0)$ optimal policy at $t = 1$ yields

$$x_1 = -\frac{(1 + \beta W)\kappa\xi}{(1 + \beta W)\kappa^2 + \lambda}\pi_0; \quad \pi_1 = \frac{\lambda}{(1 + \beta W)\kappa^2 + \lambda}\pi_0.$$

This is similar to the forward-looking model solution that reflects the trade-off between output and inflation, except that optimal policy now takes into account the cost of having inflation away from target tomorrow, through W. The loss for this case is $V(x_0, \pi_0, \rho_1^n) = W\pi_0^2/2$ since in this case the problem is the same as the one faced at $t = 2$. If $\rho_1^n < \rho_1^*(x_0, \pi_0)$ the ZLB binds, in which case

$$x_1 = \delta x_0 + \frac{\pi_0 + \rho_1^n}{\sigma}; \quad \pi_1 = \kappa\delta x_0 + \pi_0 \left(\xi + \frac{\kappa}{\sigma} \right) + \kappa\frac{\rho_1^n}{\sigma}. \tag{9}$$

The expected loss from $t = 1$ on as a function of the output gap and inflation at $t = 0$ is then given by:

$$
\begin{aligned}
L(x_0, \pi_0) &= \frac{W}{2}\pi_0^2 \int_{\rho_1^*(x_0,\pi_0)}^{+\infty} f_\rho(\rho)d\rho + \\
&\int_{-\infty}^{\rho_1^*(x_0,\pi_0)} \frac{1+\beta W}{2}\left(\kappa \delta x_0 + \pi_0\left(\xi + \frac{\kappa}{\sigma}\right) + \kappa\frac{\rho}{\sigma}\right)^2 + \frac{\lambda}{2}\left(\delta x_0 + \frac{\pi_0 + \rho}{\sigma}\right)^2 f_\rho(\rho)d\rho.
\end{aligned}
$$

This expressions reveals that the initial conditions x_0 and π_0 matter by shifting (i) the payoff from continuation in the non-ZLB states, $W\pi_0^2/2$, (ii) the payoff in the case where the ZLB binds (the second integral), and (iii) the relative likelihood of ZLB and non-ZLB states through $\rho_1^*(x_0, \pi_0)$. Since the loss function is continuous in ρ (even at $\rho_1^*(x_0, \pi_0)$), this last effect is irrelevant for welfare at the margin.

The last step is to find the optimal policy at time 0, taking into account the effect on the expected loss tomorrow:

$$
\min_{x_0} \frac{1}{2}\left((\xi\pi_{-1} + \kappa x_0 + u_0)^2 + \lambda x_0^2\right) + \beta L(x_0, \pi_0) \text{ s.t. } x_0 \le \delta x_{-1} + \frac{\rho_0^n + \pi_{-1}}{\sigma}.
$$

We use this expression to prove the following, which is analogous to Proposition 1:

Proposition 2 *For any initial condition, a mean-preserving spread in the distribution of the natural rate tomorrow ρ_1^n leads to a looser optimal policy today.*

The proof of Proposition 2 is in the appendix. Note that it incorporates the case of uncertainty regarding cost-push shocks at $t = 1$ and shows that a mean preserving spread in the cost-push shock tomorrow leads to looser policy today as well.

2.2.2 Discussion

As far as we know Proposition 2 is a new result, but its implications are similar to those of Proposition 1. As in the forward-looking case lift-off from an optimal zero interest rate should be delayed today with an increase in uncertainty about the natural rate or cost-push

shock that raises the odds of the ZLB binding tomorrow. Similarly, from an interest rate higher than the ZLB, an increase in uncertainty about the natural rate or cost-push shocks that raises uncertainty about the likelihood of being constrained by the ZLB tomorrow leads to a faster drop in the policy rate today. So the buffer stock channel and the expectations channel have very similar policy implications but for very different reasons. The expectations channel involves the possibility of being constrained by the ZLB tomorrow feeding backward to looser policy today. The buffer stock channel has looser policy today feeding forward to reduce the likelihood and severity of being at the ZLB tomorrow.

It is useful to compare the policy implications of the buffer stock channel to the argument developed in Coibion et al. (2012). That paper studies the ZLB in the context of policy reaction functions rather than optimal policy. It finds that an increase in the central bank's inflation target can reduce the likelihood and severity of policy being constrained by the ZLB. Our analysis does not require such a drastic change in monetary policy in order to improve outcomes; this is accomplished through standard interest rate policy. It is not necessary for the central bank to resort to a change to its inflation target which could damage its hard-earned credibility.[22]

2.3 Quantitative Assessment

The previous sections demonstrates how higher uncertainty about future conditions justifies a looser policy. In the current environment this implies delayed liftoff. To assess the quantitative magnitudes of these effects, and to illustrate some qualitative features of the solution, this section constructs some more realistic examples, which are solved numerically. Using parameters drawn from the literature and an estimate of the conditions prevailing today, we compare the outcomes from optimal policy (under discretion) with the well-known "Taylor rule" as specified in Taylor (1993).

---

[22]Coibion et al. (2012) do their analysis within the context of a medium-scale DSGE model with both forward and backward looking elements. It remains to be seen how important the buffer stock channel is for outcomes in a medium-scale DSGE setting.

2.3.1   Parameter values

The parameters underlying our quantitative analysis are reported in Table 1. The time period is one quarter and we take $t = 1$ to be 2015q1. The natural real rate of interest ρ_t^n is assumed to rise linearly between $t = 1$ and $t = T_0$, after which it remains constant at $\bar{\rho} = 1.75\%$. This value corresponds to the median SEP forecast for the long-run FFR (3.75%), less the FOMC's inflation target (2%). We take T_0 to be 4 years which is consistent with most SEP forecasts ending 2017 within or somewhat below the range of long run FFR in the SEP. Both models are subject to natural rate and cost-push shocks that are independent AR(1) processes with auto-correlation coefficients ρ_ε and ρ_u and innovation standard deviations σ_ε and σ_u. We use values for these parameters that are similar to the literature with the cost-push shock less persistent and less volatile than the natural rate shock.[23]

Our calculations assume that there is a date $T > T_0$ such that, after time T, the natural rate is a constant equal to $\bar{\rho}$, and there are no further shocks. This allows us to calculate the equilibrium fairly easily by backward induction under different policy rules. The precise value of T does not matter for our results, provided it is large enough. Details of the computational methods used for each model are available in the online appendix. Importantly, and in contrast to most of the literature, our numerical methods allow for uncertainty to affect policy, and to be reflected in welfare.[24]

We use the same values for the other parameters that are common to the two models. The Phillips curve slope κ is 0.02; the elasticity of substitution is $1/\sigma = 0.5$; and the discount factor is $\beta = 0.995$, all standard settings in the New Keynesian literature. The Taylor rule has weight $\phi = 1.5$ on inflation, weight $\gamma = 0.5$ on the output gap, and constant term equal to 3.75% (the "correct" long-run natural real rate 1.75% plus the 2% inflation target.)

For the backward-looking model we also need values for the coefficients on the lagged

---

[23]See for instance Adam and Billi (2007) for a calibration of a similar model, and Laubach and Williams (2003) for estimates of the volatility of the natural rate.

[24]For instance the FRB/US model used at the Reserve Board for policy simulations cannot address uncertainty systematically. Hamilton et al. (2015) use this model to assess the implications of assuming different values of the natural rate within their wide range of estimates. However by using FRB/US they are unable to address the impact of uncertainty *per se* on optimal policy.

Table 1: Parameter Values

Parameter	Description	Value
β	Discount factor	0.995
κ	Slope of Phillips Curve	0.02
σ	Inverse elasticity of substitution	2
δ	Backward-looking IS curve coef.	0.75
ξ	Backward-looking Phillips curve coef.	0.95
σ_ε	Std. dev. natural rate innovation	0.3
σ_u	Std. dev. of cost-push innovation	0.15
ρ_ε	Serial correlation of natural rate	0.92
ρ_u	Serial correlation of cost-push	0.3
λ	Weight on output stabilization	0.25
π^*	Steady-state inflation (annualized)	2
$\overline{\rho}$	Terminal natural rate (annualized)	1.75
T_0	Quarters to reach terminal natural rate	16
ρ_0^n	Value of natural rate at time 1	-0.5
ϕ	Taylor rule coefficient on inflation	1.5
γ	Taylor rule coefficient on output gap	0.5
x_0	Initial condition for the output gap	-1.5
π_0	Initial condition for inflation	1.3

Note: Values of standard deviations, inflation, the output gap, and the natural rate are shown in percentage points.

terms in (6) and (7) as well as initial conditions for the output gap and inflation. The coefficient on lagged inflation, $\xi = 0.95$, is based on a simple instrumental variables estimation of (6). The coefficient on lagged output in the backward-looking model's IS curve is $\delta = 0.75$, in order to generate significant persistence in the output gap. We assume an initial inflation rate of 1.3%, the latest reading for core PCE inflation as we write this paper, and an initial output gap $x_0 = -1.5\%$. There is obviously substantial uncertainty about the level of the output gap, but a simple or naive calculation using the 2014q4 unemployment rate (5.7%), an estimate of the natural rate of unemployment rate (5.0%) and Okun's law yields -1.5%. The CBO reports the output gap to be -2.5% as of this writing.

2.3.2 Forward-looking model results

Figure 1 displays the path of interest rates, inflation and the output gap under two policies: optimal policy under discretion, and the Taylor rule. These are the "baseline" paths, *i.e.* those that arise if the realized shocks to the natural rate and to the cost-push variable are identically zero, though optimal policy is set as if they were uncertain. The interest rate panel also displays the nominal natural rate for comparison. As discussed in Section 2.1, if there was no uncertainty, setting the interest rate equal to this nominal natural rate would implement the optimal outcome with zero output gap and inflation gaps ($x = 0$ and $\pi = \pi^*$). Hence, the difference between the path for the nominal natural rate and optimal policy demonstrates the extent to which uncertainty affects optimal policy.

Figure 1: Lift-off in the Forward-looking Model
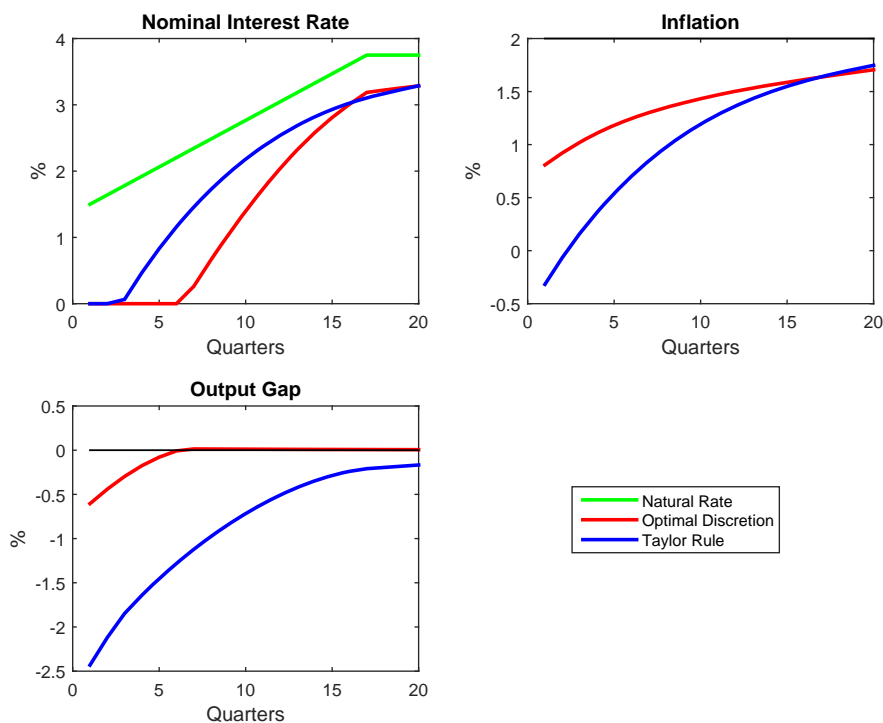
In the forward-looking model, the initial output and inflation gaps are negative, and converge back to zero over time. These gaps are not initial conditions but are endogenous to

uncertainty and to the policy response. They arise because agents anticipate the possibility of future negative shocks that the central bank may be powerless to offset due to the ZLB. As the natural rate rises, this risk diminishes and optimal policy lifts off, and eventually converges to the natural rate. In this simulation, lift off occurs in period 6, *i.e.* 2016q2. The delayed lift-off relative to the path of the natural rate is entirely the consequence of the uncertainty faced by the central bank. By comparison, the Taylor rule, which does not take into account this uncertainty, implies faster liftoff, in the third period (2015q3). This faster liftoff (as well as expectations of reactions to possible shocks) leads to lower expected output and inflation that translate into lower output and inflation today.[25]

Table 2 compares some outcomes under these two policies, taking into account the uncertainty by simulating 50,000 paths of the model with different realizations of the shocks. Clearly, optimal policy under discretion achieves a much lower loss, because it recognizes the risk of negative shocks driving the economy back to the ZLB.[26] Time-to-liftoff can be a difficult statistic to interpret, however, because a reaction function that depends on output and inflation might have a late liftoff because it is too restrictive and thus generates low output and inflation. For this reason, we find it useful to report the *typical* conditions under which liftoff occurs, *i.e.* the median (across simulations) of inflation and output gap at liftoff. We find that the Taylor rule typically lifts off when the output gap is -1.5% and inflation is 1%. But optimal policy lifts off when the output gap is close to zero and inflation is 1.5%.

When comparing policies it is also important to balance the risks associated with each. The loss function is a relevant summary statistic, but we find it helpful to complement it with simpler summary statistics. We report for each policy the median (across simulations) of the maximum inflation over the next 5 years; and similarly the median (across simulations) of the lowest output gap over the next 5 years. Optimal policy reduces the risk of very low output gaps substantially, with the median minimum going from -4.5% to -1.7%. But,

---

[25]Indeed, the output gap and inflation at $t = 1$ are lower than what we observe currently. This might reflect the effect of some shocks at $t = 1$ that we do not model, or might simply reflect that the Fed is not following the Taylor rule.

[26]This result does not have to hold by definition. The Taylor rule provides commitment which may lead to more favorable outcomes, for instance if cost-push shocks are volatile and persistent.

Table 2: Forward-looking Simulation

Statistic	Optimal Policy	Taylor Rule
Loss	0.12	0.76
Median time at liftoff	6	3
Median π at liftoff	1.50	1.05
Median x at liftoff	0.01	-1.53
Median $\max(\pi)$	3.57	4.55
Median $\min(x)$	-1.74	-4.47

perhaps surprisingly, optimal policy also reduces the risk of a very high inflation, and reduces the typical maximum inflation from 4.6% to 3.6%. This is because optimal policy responds quite strongly to shocks that create inflation.

2.3.3 Backward-looking model results

Figure 2 is the analogue of Figure 1 for the backward-looking model. As in the case of the forward-looking model, we see that optimal policy under discretion is significantly looser than the Taylor rule, which lifts off right away. To illustrate that higher uncertainty leads to delayed liftoff, we also solve the optimal policy when the standard deviation of shocks is 50% larger. Liftoff is even more delayed under these circumstances (even though our simulation assumes that no shocks are actually realized).

Optimal policy thus achieves a faster return to low gaps, and builds up a buffer during this transition so as to guard against bad shock realizations. The output gap and inflation optimally overshoot their targets during the transition. Table 3 is analogous to Table 2. In the backward-looking model as well, optimal policy significantly outperforms the Taylor rule according to the expected loss. Median liftoff occurs in 2016q3, but is state-contingent of course, and typically occurs with inflation essentially at target (2.0%) and the output gap positive (0.2%). By comparison, the Taylor rule lifts off with inflation and the output gap well below target.[27]

---

[27]Indeed, the Taylor rule lifts off even earlier than these statistics suggest, since we start our computation at time 1 and it has already lifted off.

Figure 2: Lift-off in the Backward-looking Model
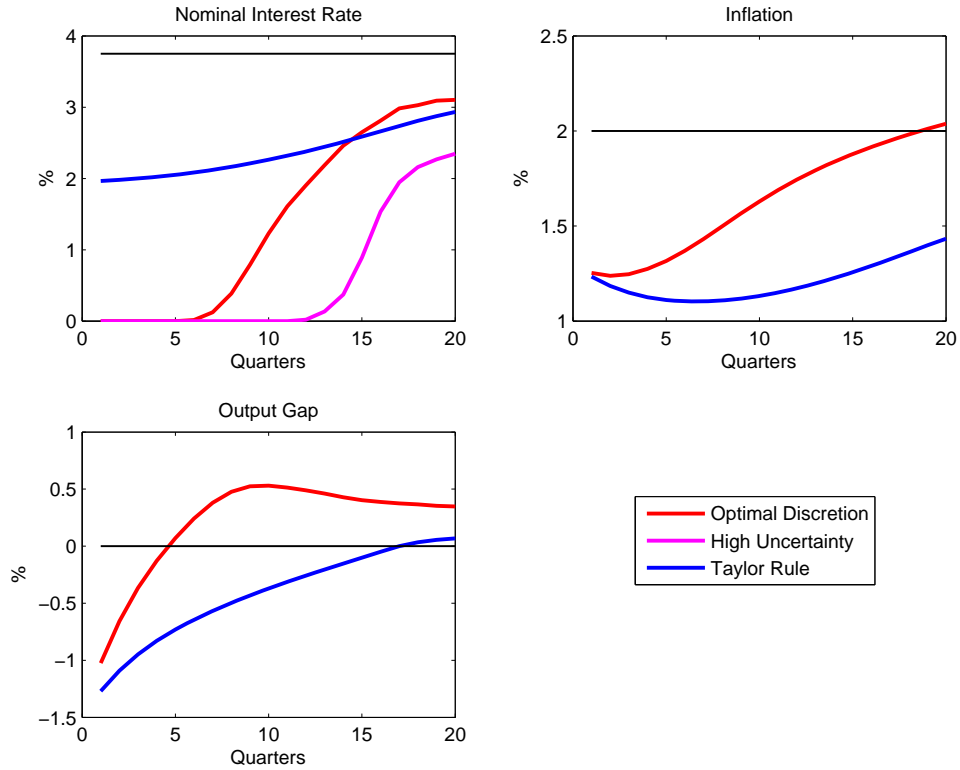
Table 3: Backward-looking Simulation

Statistic	Optimal Policy	Taylor Rule
Loss	0.16	0.36
Median time at liftoff	7	1
Median π at liftoff	1.96	1.23
Median x at liftoff	0.18	-1.27
Median $\max(\pi)$	4.30	4.08
Median $\min(x)$	-1.03	-1.46

One difference with the previous model is that the Taylor rule does slightly better at reducing peak inflation, with the typical maximum inflation reduced from 4.3% under optimal policy to 4.1%. As in the previous case however, it is worst at preventing bad outcomes, with the minimum output gap going from -1.0% to -1.5%. Overall, the backward-looking model, despite its very different structure, generates outcomes that are similar to the forward-looking model.

Figure 3: Large Cost-push Shock in the Backward-looking Model
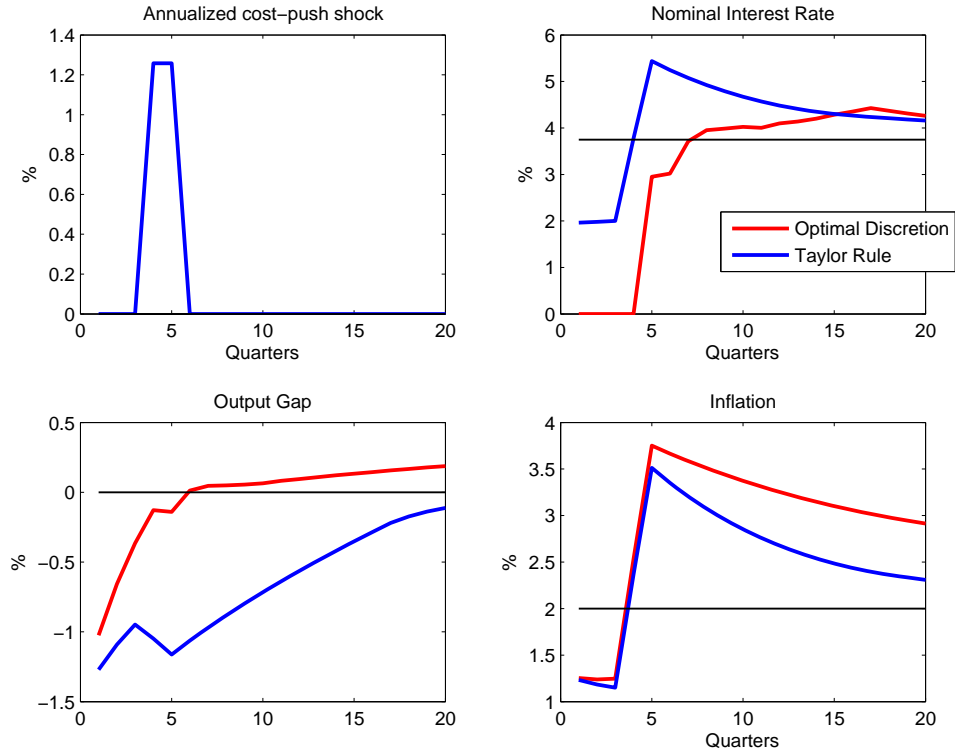
We conclude by illustrating one of the risks the optimal policy is able to address, namely the possibility that a shock will drive up inflation before the baseline lift-off. Figure 3 depicts a particular simulation where there are two consecutive large positive cost-push shocks that hit the economy before the baseline date of lift-off. The shocks trigger a large increase in interest rates, enough so that the inflation response is only mildly larger than when the economy follows the Taylor rule. Similarly, if a large positive shock to the natural rate were to hit the economy, higher interest rates would follow, and in this case would stabilize both inflation and output. The simple logic is that staying at zero longer under the baseline scenario does not impair the ability of monetary policy to respond to future contingencies.[28] We obtain similar results with the forward looking model.

---

[28]One potential criticism of the optimal policy is that it may entail large movements in interest rates. Our policy calculation does not give any weight to "interest-rate smoothing." It is difficult to rationalize interest-rate smoothing however; indeed some authors argue that the smoothness of interest rates in the data is due to learning rather than a desire to smooth interest rates *per se* (Sack (2000) and Rudebusch (2001)).

3 Historical Precedents for Risk Management

The FOMC's historical policy record provides many examples of how risk management considerations likely have influenced monetary policy decisions. FOMC minutes and other Federal Reserve communications reveal a number of episodes when the FOMC indicated that it took a wait-and-see approach to taking further actions or muted a FFR move due to its uncertainty over the course of the economy or the extent to which the full force of early policy moves had yet shown through to economic activity and inflation. The record also indicates several instances when the Committee said its policy stance was taken in part as insurance against undesirable outcomes; during these times, the FOMC also usually noted reasons why the potential costs of a policy overreaction likely were modest as compared with the scenario it was insuring against. And there are a few instances when the Committee appears to be reacting to head off a potential change in dynamics that might accelerate the economy into a serious recession.

Two episodes are particularly revealing. The first is the hesitancy of the Committee to raise rates in 1997 and 1998 to counter inflationary threats because of the uncertainty generated by the Asian financial crisis and subsequent rate cuts following the Russian default. The second is the loosening of policy over 2000 and 2001, when uncertainty over the degree to which growth was slowing and the desire to insure against downside risks appeared to influence policy. Furthermore, later in the period, the Committee's aggressive actions also seemed to be influenced by attention to the risks associated with the ZLB on interest rates.

Of course, not all references to risk management involved reactions to uncertainty or insurance-based rationales for policy. For example, at times the FOMC faced conflicting policy prescriptions for achieving its dual mandate goals for output and inflation. Here, the Committee generally hoped to set policy to better align the risks to the projected deviations from the two targets—an interesting balancing act, though not necessarily one in which higher moments of the distribution of shocks or potential nonlinearities in economic dynamics had a meaningful influence of policy decisions.

The ZLB was not a relevant consideration for most of the historical record we consider as well as for the empirical work we do later in Section 4. Accordingly, this section begins by briefly reviewing some theoretical rationales for taking a risk management approach to policy away from the ZLB that may help shed light on some Federal Reserve actions. We will then describe the two episodes we find particularly revealing about the use of risk management in setting rates. The section concludes with two approaches to quantifying the role of risk management as it is described in the FOMC minutes.

3.1   Rationales for risk management away from the ZLB

Policymakers have long-emphasized the importance of uncertainty in their decision- making. As Greenspan (2004) put it: "(t)he Federal Reserve's experiences over the past two decades make it clear that uncertainty is not just a pervasive feature of the monetary policy landscape; it is the defining characteristic of that landscape." This sentiment seems at odds with a wide class of models (namely linear-quadratic models) in which optimal policy involves adjusting the interest rate in response to the mean of the distribution of shocks and information on higher moments is irrelevant (away from the ZLB). What kinds of factors cause departures from such conditions and justify the risk-management approach?

Nonlinearities in economic dynamics are one natural motivation. These can occur in both the IS and Phillips curves. For example, suppose recessions are episodes when self-reinforcing dynamics amplify the effects of downside shocks. This could be modeled as a dependence of output on lagged output, as in the backward-looking model studied above, but this dependence may be concave rather than linear. Intuitively, negative shocks have a more dramatic effect on reducing future output than positive shocks have on increasing it, and so the greater uncertainty, the more optimal policy will be looser *ex-ante* to guard against the more detrimental outcomes. Alternatively, suppose the Phillips curve is convex, perhaps owing to wage rigidities at low inflation. Intuitively, a positive shock to the output gap leads to a significant increase of inflation above target while a negative shock does not lead to much of an inflation decline; the larger the possible spread of these shocks, the greater

the relative odds of experiencing a bad inflation outcome. Optimal policy will want to guard against this, leading to a tightening bias.[29]

Relaxing the assumption of a quadratic loss function is perhaps the simplest way to generate a rationale for risk-management. The quadratic loss function is justified as being a local approximation to consumer welfare (Woodford (2003)). However, it might not be a very good approximation during times when large shocks drive the economy far from the underlying trend or the quadratic loss function might simply be an inadequate approximation of the way the FOMC actually behaves. Several authors have considered loss functions with asymmetries, *e.g.* Surico (2007), Kilian and Manganelli (2008), Dolado, María-Dolores, and Ruge-Murcia (2004). For instance, the latter paper shows that if the policymaker is less averse to output running above potential than below it, then the optimal policy rule can involve nonlinear terms in the output and inflation gaps. The relevance of higher moments of the distribution for shocks for the policy decision is an obvious by-product of these nonlinearities. To the degree that optimal policy depends on the second and third moments of the shock distributions, both uncertainty and asymmetry may enter the policymaking calculus.

The risk management approach also can be found in the large literature on how optimal monetary policy should adjust for uncertainty about the true model of the economy. Brainard (1967) derived the important result that parameter uncertainty over the effects of policy should lead to additional caution and smaller policy responses to deviations from target, a principle that is often called "gradualism." This principle has had considerable influence on policymakers, for instance Blinder (1998) or Williams (2013). There also is a more recent literature that incorporates concern about model miss-specification into optimal monetary policy analysis. Sometimes this is along the lines of the robust control analysis of Hansen and Sargent (2008), which often has been interpreted to mean that model uncertainty should generate more aggressive policy. However, as explained by Barlevy (2011), both Brainard-style gradualism and robust-control aggressive policy results depend on the specifics of the underlying models of the economy, and reasonable alternatives can overturn the baseline

---

[29]The fact that a convex Phillips curve can lead to a role for risk-management has been discussed by Laxton, Rose, and Tambakis (1999) and Dolado, María-Dolores, and Naveira (2005).

results. Hence, the effect of parameter and model uncertainties are themselves uncertain. Nonetheless, these analyses often indicate that higher moments of the distribution of shocks can influence the setting of optimal policy.

With this as background, we now turn to the two historical examples we think are particularly insightful for how the Federal Reserve may have implemented the risk-management approach to policy.

3.2 1997–1998

1997 was a good year for the U.S. economy: real GDP increased 3-3/4 percent, the unemployment rate fell to 4.7 percent—about 3/4 percentage point below the Board of Governors staff's estimate of the natural rate–and core CPI inflation was 2-1/4 percent.[30] But with growth solid and labor markets tight, the FOMC clearly was concerned about a buildup in inflationary pressures. As noted in the Federal Reserve's February 1998 Monetary Policy Report:

> The circumstances that prevailed through most of 1997 required that the Federal Reserve remain especially attentive to the risk of a pickup in inflation. Labor markets were already tight when the year began, and nominal wages had started to rise faster than previously. Persistent strength in demand over the year led to economic growth in excess of the expansion of the economy's potential, intensifying the pressures on labor supplies.

Indeed, over much of the period between early 1997 and mid-1998, the FOMC directive maintained a bias indicating that it was more likely to raise rates to battle inflationary pressures than it was to lower them. Nonetheless, the FOMC left the FFR unchanged at 5.5 percent from March 1997 until September 1998. Why did it do so?

Certainly the inaction in large part reflected the forecast for economic growth to moderate to a more sustainable pace as well as the fact that actual inflation had remained contained

---

[30]The GDP figure refers to the BEA's third estimate for the year released in March 1998.

despite tight labor market conditions.[31] But, in addition, on numerous occasions heightened uncertainty over the outlook for growth and inflation apparently reinforced the decision to refrain from raising rates. The following quote from the July FOMC 1997 minutes is a revealing example:

> An unchanged policy seemed appropriate with inflation still quiescent and business activity projected to settle into a pattern of moderate growth broadly consistent with the economy's long-run output potential. While the members assessed risks surrounding such a forecast as decidedly tilted to the upside, the slowing of the expansion should keep resource utilization from rising substantially further, and this outlook together with the absence of significant early signs of rising inflationary pressures suggested the desirability of a cautious "wait and see" policy stance at this point. In the current uncertain environment, this would afford the Committee an opportunity to gauge the momentum of the expansion and the related degree of pressure on resources and prices.

Furthermore, the Committee did not see high costs to "waiting and seeing." They thought any increase in inflation would be slow, and that if needed a limited tightening on top of the current 5.5 percent funds would be sufficient to reign in any emerging price pressures. This is seen in the following quote from the same meeting:

> The risks of waiting appeared to be limited, given that the evidence at hand did not point to a step-up in inflation despite low unemployment and that the current stance of monetary policy did not seem to be overly accommodative ... In these circumstances, any tendency for price pressures to mount was likely to emerge only gradually and to be reversible through a relatively limited policy adjustment.

Thus, it appears that uncertainty and associated risk management considerations supported the Committee decision to leave policy on hold.

---

[31]Based on the FFR remaining at 5.5 percent, the August 2008 Greenbook projected GDP growth to slow from 2.9 percent in 1998 to 1.7 percent in 1999. The unemployment rate was projected to rise to 5.1 percent by the end of 1999 and core CPI inflation was projected to edge down to 2.1 percent. Note that core PCE inflation was much lower than core CPI inflation at this time – it was projected at 1.3 percent in 1998 and 1.5 percent in 1999. However, the FOMC had not yet officially adopted the PCE price index as its preferred inflation measure, nor had it set an official inflation target.

Of course, the potential fallout of the Asian financial crisis on the U.S. economy was a major factor underlying the uncertainty about the outlook. The baseline scenario was that the associated weakening in demand from abroad and a stronger dollar would be enough to keep U.S. inflationary pressures in check but not be strong enough to cause inflation or employment to fall too low. As Chairman Greenspan noted in his February 1998 Humphrey-Hawkins testimony to Congress, there were substantial risks to this outlook, with the delicate balance dictating unchanged policy:

> However, we cannot rule out two other, more worrisome possibilities. On the one hand, should the momentum to domestic spending not be offset significantly by Asian or other developments, the U.S. economy would be on a track along which spending could press too strongly against available resources to be consistent with contained inflation. On the other, we also need to be alert to the possibility that the forces from Asia might damp activity and prices by more than is desirable by exerting a particularly forceful drag on the volume of net exports and the prices of imports. When confronted at the beginning of this month with these, for the moment, finely balanced, though powerful forces, the members of the Federal Open Market Committee decided that monetary policy should most appropriately be kept on hold.

Indeed, by late in the summer of 1998, this balance had changed, as the strains following the Russian default weakened the outlook for foreign growth and tightened financial conditions in the U.S. The Committee was concerned about the direct implications of these developments on U.S. financial markets – which were already evident in the data – as well as for the real economy, which were still just a prediction. The staff forecast prepared for the September FOMC meeting reduced the projection for growth in 1999 by about 1/2 percentage point (to 1-1/4 percent), a forecast predicated on a 75 basis point reduction in the FFR spread out over three quarters. Such a forecast was not a disaster—indeed, at 5.1 percent, the unemployment rate projected for the end of 1999 was still below the Board Staff's estimate of its natural rate inflation. Nonetheless, the FOMC moved much faster than assumed in the staff's forecast, lowering rates 25 basis points at its September and November meetings as well at an intermeeting cut in October. According to the FOMC minutes, the rate cuts

were made in part as insurance against a worsening of financial conditions and weakening activity. As they noted in September:

> ... such an action was desirable to cushion the likely adverse consequences on future domestic economic activity of the global financial turmoil that had weakened foreign economies and of the tighter conditions in financial markets in the United States that had resulted in part from that turmoil. At a time of abnormally high volatility and very substantial uncertainty, it was impossible to predict how financial conditions in the United States would evolve.... In any event, an easing policy action at this point could provide added insurance against the risk of a further worsening in financial conditions and a related curtailment in the availability of credit to many borrowers.

While the references to insurance are clear, the case also can be made that these policy moves were in large part made to realign the misses in the expected paths for growth and inflation from their policy goals. Over this time the prescriptions to address the risks to the FOMC's dual mandate policy goals were in conflict—risks to achieving the inflation mandate called for higher interest rates while risks to achieving the maximum employment mandate called for lower rates.[32] As the above quote from Chairman Greenspan indicated, in 1997 the Committee thought that a 5-1/2 percent FFR kept these risks in balance. But as the odds of economic weakness increased, the Committee cut rates to bring the risks to the two goals back into balance. As Chairman Greenspan indicated in his February 1999 Monetary Policy Testimony:

> To cushion the domestic economy from the impact of the increasing weakness in foreign economies and the less accommodative conditions in U.S. financial markets, the FOMC, beginning in late September, undertook three policy easings. By mid-November, the FOMC had reduced the federal funds rate from 5-1/2

---

[32]To quote the February 1999 Monetary Policy Report: "Monetary policy in 1998 needed to balance two major risks to the economic expansion. On the one hand, with the domestic economy displaying considerable momentum and labor markets tight, the Federal Open Market Committee (FOMC) was concerned about the possible emergence of imbalances that would lead to higher inflation and thereby, eventually, put the sustainability of the expansion at risk. On the other hand, troubles in many foreign economies and resulting financial turmoil both abroad and at home seemed, at times, to raise the risk of an excessive weakening of aggregate demand."

percent to 4-3/4 percent. These actions were taken to rebalance the risks to the outlook, and, in the event, the markets have recovered appreciably.

So were the late 1998 rate moves a balancing of forecast probabilities, insurance against a downside skew in possible outcomes, or some combination of both? There is no easy answer. This motivates our econometric work in Section 4 that seeks to disentangle the normal response of policy to expected outcomes from uncertainty and other related factors that may have influenced the policy decision.

In the end, the economy weathered the fallout from the Russian default well. In June 1999, the staff forecast projected the unemployment rate to end the year at 4.1 percent and that core CPI inflation would rise to 2.5 percent by 2000.[33] Against this backdrop, the FOMC decided to increase the FFR to 5 percent. In the event, the staff forecast underestimated the strength of the economy and underlying inflationary pressures, and the FOMC ended up executing a series of rate hikes that eventually brought the FFR up to 6.5 percent by May of 2000.

3.3   2000–2001

At the time of the June 2000 FOMC meeting, the unemployment rate stood at 4 percent and core PCE inflation, which the Committee was now using as its main measure of consumer price inflation, was running at about 1-3/4 percent, up from 1-1/2 percent in 1999. The staff forecast growth would moderate to a rate near or a little below potential but that unemployment would remain near its current level and that inflation would rise to 2.3 percent in 2001—and this forecast was predicated on another 75 basis points tightening that would bring the FFR to 7-1/4 percent by the end of 2000. Despite this outlook, the FOMC decided to leave rates unchanged. What drove this pause? It seems likely to us that risk management was an important consideration.

In particular, the FOMC appeared to want to see how uncertainty over the outlook

---

[33]This forecast was based on an assumption of the FFR gradually moving up to 5-1/4 percent by the first quarter of 2000.

would play out. First, the incoming data and anecdotal reports from Committee members' business contacts pointed to a slowdown in growth, but the degree of the slowing was not clear. Second, rates had risen substantially over the past year, and given the lags from policy changes to economic activity, it was unlikely that the full effects of the hikes had yet been felt. Given the relatively high level of the FFR and the slowdown in growth that appeared in train, the Committee seemed wary of over tightening. Third, despite the staff forecast, the FOMC apparently considered the costs of waiting in terms of inflation risks to be small. Accordingly, they thought it better to put a rate increase on hold and see how the economy developed. The June 2000 minutes contain a good deal of commentary supporting this interpretation:[34]

> The increasing though still tentative indications of some slowing in aggregate demand, together with the likelihood that the earlier policy tightening actions had not yet exerted their full retarding effects on spending, were key factors in this decision. The uncertainties surrounding the outlook for the economy, notably the extent and duration of the recent moderation in spending and the effects of the appreciable tightening over the past year ... reinforced the argument for leaving the stance of policy unchanged at this meeting and weighting incoming data carefully. . . .Members generally saw little risk in deferring any further policy tightening move, particularly since the possibility that underlying inflation would worsen appreciably seemed remote under prevailing circumstances. Among other factors, inflation expectations had been remarkably stable despite rising energy prices, and real interest rates were already relatively elevated.

Moving through the second half of 2000, it became increasingly evident that growth had slowed to a pace somewhat below trend and may in fact have been poised for even more pronounced weakness. Furthermore, inflation was moving up at a slower pace than the staff

---

[34]This was not the first time the Committee had invoked such arguments during the tightening cycle. In October 1999 the FOMC left rates unchanged in part over uncertainty over the economic outlook. And in the February and March 2000 meetings they opted for small 25 basis point cuts because of uncertainty. As stated in the July 2000 Monetary Policy Report to Congress regarding the smaller moves in February and March: "The FOMC considered larger policy moves at its first two meetings of 2000 but concluded that significant uncertainty about the outlook for the expansion of aggregate demand in relation to that of aggregate supply, including the timing and strength of the economy's response to earlier monetary policy tightenings, warranted a more limited policy action."

had projected in June. In response, the Committee held the FFR at 6.5 percent through the end of 2000. But the data around the turn of the year proved to be weaker than the Committee had anticipated. In a conference call on January 3, 2001, the FOMC cut the FFR to 6 percent and lowered it again to 5-1/2 percent at the end-of-month FOMC meeting.[35] In justifying the aggressive ease, the Committee stated:

> Such a policy move in conjunction with the 50 basis point reduction in early January would represent a relatively aggressive policy adjustment in a short period of time, but the members agreed on its desirability in light of the rapid weakening in the economic expansion in recent months and associated deterioration in business and consumer confidence. The extent and duration of the current economic correction remained uncertain, but the stimulus ... would help guard against cumulative weakness in economic activity and would support the positive factors that seemed likely to promote recovery later in the year ... In current circumstances, members saw little inflation risk in such a "front-loaded" easing policy, given the reduced pressures on resources stemming from the sluggish performance of the economy and relatively subdued expectations of inflation.

According to this quote, not only was the actual weakening in activity an important consideration in the policy decision, but uncertainty over the extent of the downturn – and the possibility that it might turn into an outright recession – seemed to spur the Committee to make a large move. The "help guard against cumulative weakness" and "front-loaded" language could be read as the Committee taking out some additional insurance against the possibility that the weakening activity would snowball into a recession. Indeed this could have reflected a concern about the kinds of non-linear output dynamics or perhaps non-quadratic losses associated with a larger recession that we discuss in Section 3.1.

The FOMC steadily brought the FFR down further over the course of 2001 against the backdrop of weakening activity. Still, though the economy seemed to be skirting a recession. Then the tragic events of September 11 occurred. There was, of course, huge uncertainty

---

[35] At that time the Board staff was forecasting that growth would stagnate in the first half of the year, but that the economy would avoid an outright recession even with the FFR at 5.75 percent. Core PCE inflation was projected to rise modestly to a little under 2.0 percent.

over how international developments, logistics disruptions, and the sentiment of households, businesses, and financial markets would affect spending and production. By November the Board staff was forecasting a modest recession: Growth in the second half of 2001 was projected to decline 1-1/2 percent at an annual rate and rise at just a 1-1/4 percent rate in the first half of 2002. By the end of 2002 the unemployment rate was projected to rise to 6.1 percent and core PCE inflation was projected to be 1-1/2 percent. These forecasts were predicated on the FFR remaining flat at 2-1/4 percent.

The FOMC, however, was worried about something more serious than the shallow recession forecast by the Staff. Furthermore, a new risk came to light, namely the chance that disinflationary pressures might emerge, that, once established, would be more difficult to fight with the FFR already low. In response, the Committee cut the FFR 50 basis points in a conference call on September 17 and again at their regular meetings in October and November. As earlier in the year, they preferred to act aggressively. As noted in the minutes from the November 2001 FOMC meeting:

> ...members stressed the absence of evidence that the economy was beginning to stabilize and some commented that indications of economic weakness had in fact intensified. Moreover, it was likely in the view of these members that core inflation, which was already modest, would decelerate further. In these circumstances insufficient monetary policy stimulus would risk a more extended contraction of the economy and possibly even downward pressures on prices that could be difficult to counter with the current federal funds rate already quite low. Should the economy display unanticipated strength in the near term, the emerging need for a tightening action would be a highly welcome development that could be readily accommodated in a timely manner to forestall any potential pickup in inflation.

This passage suggests that the large cuts were not only aimed at preventing the economy from falling into a serious recession with deflationary pressures, but that the Committee was also concerned that such an outcome "could be difficult to counter with the current funds rate already quite low." Accordingly, the aggressive policy moves could in part also have

reflected insurance against the future possibility of running into the ZLB, precisely the policy scenario and optimal policy prescription described in Section 2.

3.4   Quantifying References to Uncertainty and Insurance in FOMC Minutes

Clearly, the minutes contain many references to the Committee noting that uncertain economic conditions influenced their policy decision and times when insurance was cited as a reason to alter the stance of policy one way or the other. But did these references actually lead to deviations from the canonical policy response based on simple point forecasts for the outlook? In this section we take a systematic approach to quantifying these considerations into variables that can be used in empirical work to address this question.

In the spirit of the narrative approach pioneered by Romer and Romer (1989), we built judgmental indicators based on our reading of the minutes over the period 1993 to 2008.[36] We concentrated on the paragraphs that describe the Committee's rationale for its policy decision, reading these passages for references to when insurance considerations or uncertainty over the economic environment or the efficacy of current or past policy moves appeared closely linked to the FOMC's decision. Other portions of the minutes were excluded from our analysis– for example, the parts that cover staff and participants' views of current and prospective economic and financial developments – in order to better isolate arguments that directly influenced the meeting's policy decision from more general discussions of unusual data or normal forecast uncertainty.

We constructed two separate indicator variables – one for uncertainty (hUnc) and one for insurance (hIns) – "h" indicates these are "human" coded indicators. The uncertainty variable was coded to plus one if we judged that the Committee positioned the FFR higher than it otherwise would due to uncertainty. We coded a minus one if it appeared that uncertainty led the FOMC to put rates lower than they otherwise would be. If uncertainty did not appear to be an important factor influencing the policy decision, we coded the indicator

---

[36]The start date is predicated on the fact that FOMC minutes prior to 1993 provide little information about the rationale for policy decisions. See Hansen and McMahon (2014) for a detailed discussion of this change in Federal Reserve communications.

as zero. We similarly coded the insurance variable by identifying when the minutes cited insurance against some adverse outcome as an important consideration in the Committee's decision.[37]

As an example of our coding, consider the June 2000 pause in rate hikes discussed above. Though they generally thought policy had to tighten, the Committee was uncertain about how much growth was slowing and the degree to which their past tightening actions had yet shown through to economic activity. Accordingly, the FOMC decided to wait and assess further developments before taking additional policy action. This is clear from the sections of the minutes highlighted in italics:

> The increasing though still tentative indications of some slowing in aggregate demand, together with the likelihood that the earlier policy tightening actions had not yet exerted their full retarding effects on spending, were key factors in this decision. *The uncertainties surrounding the outlook for the economy, notably the extent and duration of the recent moderation in spending and the effects of the appreciable tightening over the past year*, including the 1/2 percentage point increase in the intended federal funds rate at the May meeting, *reinforced the argument for leaving the stance of policy unchanged at this meeting and weighting incoming data carefully.*

We coded this meeting as a minus one for our uncertainty measure – rates were lower because uncertainty over the economic outlook and the effects of past policy moves appear to have been an important factor in the Committee deciding not raising rates when they otherwise would have.

However, we did not code all mentions of uncertainty as a one or minus one. For example, in March 1998 – a meeting when the FOMC did not change rates despite some concern over higher inflation – the Committee did refer to uncertainties over the economic outlook and say that it could wait for further developments before tightening. The FOMC had held the FFR

---

[37]A value of one for either variable could reflect the Committee raising rates by more or lowering rates by less than they would have if they ignored uncertainty or insurance or a decision to keep the FFR at its current level when a forecast-only call would have been to lower rates. Similarly, a value of minus one could occur if the FOMC either lowered rates them by more or increased them less than they otherwise would or if the Committee left rates unchanged when they otherwise would have raised them.

flat at 5.5 percent for about a year, and so was not obviously in the midst of a tightening cycle; the baseline forecast articulated in the policy paragraphs seems consistent with the FFR setting at the time; and the commentary over the need to tighten was in reference to an indefinite point in the future as opposed to the current or subsequent FOMC meeting. So, in our judgment, uncertainty did not appear to be a very important factor holding back a rate increase at this meeting and we coded this date as a zero. Quoting the minutes (again, with our emphasis added):

> The members agreed that should the strength of the economic expansion and the firming of labor markets persist, *policy tightening likely would be needed at some point to head off imbalances that over time* would undermine the expansion in economic activity. *Most saw little urgency to tighten policy at this meeting, however. The economy might well continue to accommodate relatively robust economic growth and a high level of resource use for an extended period without a rise in inflation* . . . On balance, *in light of the uncertainties in the outlook and given that a variety of special factors would continue to contain inflation for a time, the Committee could await further developments bearing on the strength of inflationary pressures without incurring a significant risk* that disruptive policy actions would be needed later in response to an upturn in inflation and inflation expectations.

Of course, coding the minutes in this way is inherently subjective and there is no definitive way to judge the accuracy of the decisions we made. So we also constructed objective measures of how often references to uncertainty or insurance appeared in the policy paragraphs of the minutes. In particular we constructed conditional measures which count the percentage of sentences containing words related to uncertainty or insurance in conjunction with references to economic activity or inflation. The words we used to capture uncertainty are "uncertainty," "uncertain," "uncertainties," "question" and "questions." To capture insurance we used "insurance," "ensure," "assurance" and "risk management." The conditioning words for inflation were "inflation," prices," "deflation," "disinflation," "cost" and "costs." To condition on activity we used "activity," "growth," "slack," "resource," "resources," "labor" and "employment." We combined the counts for uncertainty and insurance into the

two variables mUnc and mIns – "m" indicates these are "machine" coded indicators

Figure 4: Uncertainty Sentence Count and Indicator
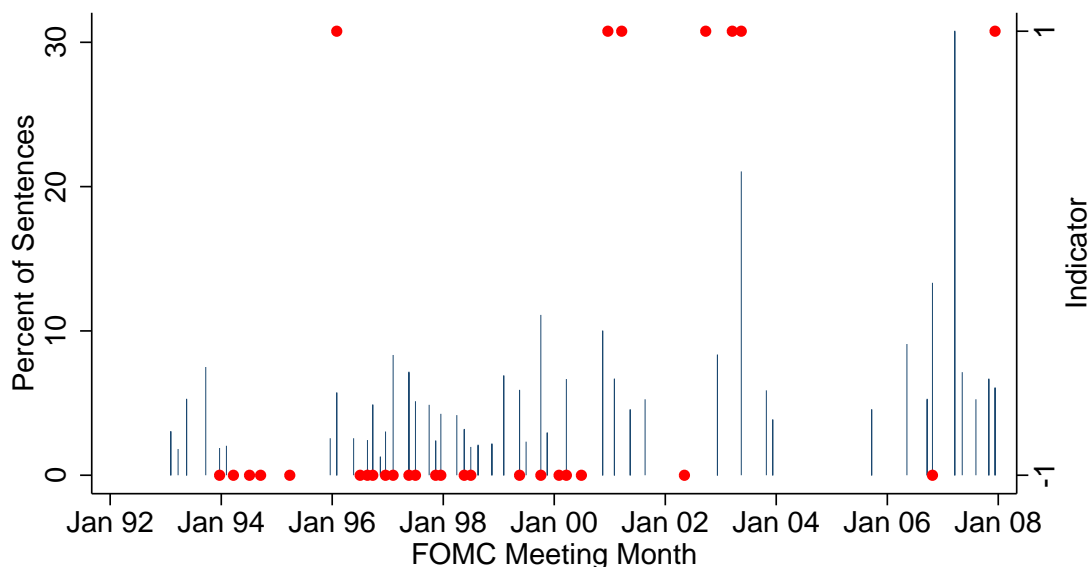
Figures 4 and 5 show plots of these uncertainty and insurance measures. Non-zero values of the indicator variables are indicated by red circles and the blue bars indicate the sentence counts. Not surprisingly, dealing with uncertainty is a regular feature of monetary policy decision-making. The uncertainty indicator "turns on" in 31 out of the 128 meetings between 1993 and 2008. Indications that insurance was a factor in shading policy are not as common, but still show up 14 times in the indicator. Most of the time – 24 for uncertainty and 11 for insurance – it appears that rates were set lower than otherwise would have been to account for these factors.

The sentence counts and indicator variables do not line up perfectly. Sometimes the indicator variables are reflected in the sentence counts but sometimes they are not. There are also meetings where the sentence counts are positive but we did not judge them to indicate

Figure 5: Insurance Sentence Count and Indicator
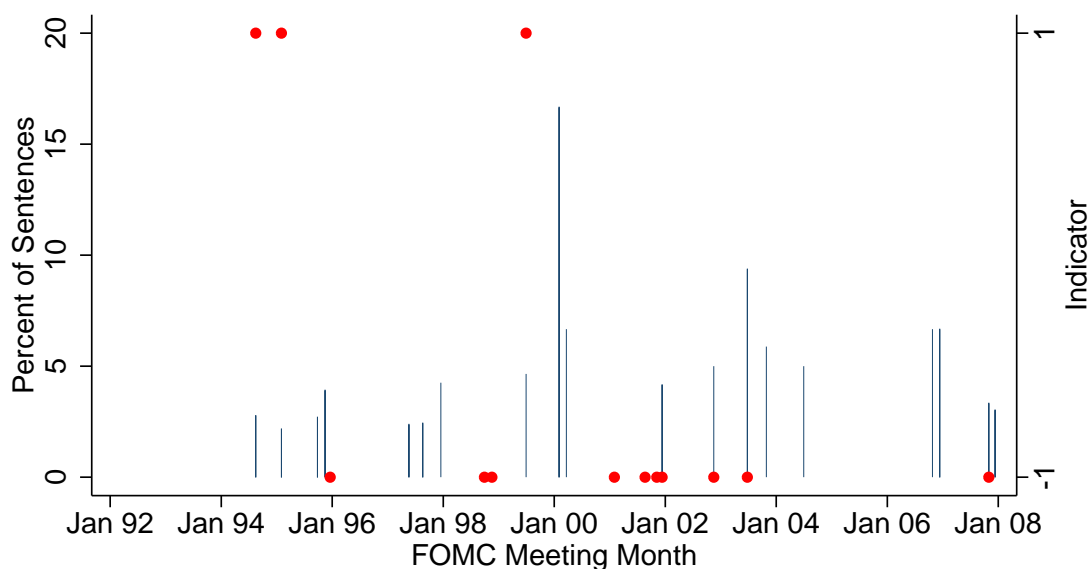
that rates were set differently than they "normally" would. For example, in March of 2007, our judgmental measure does not code uncertainty as being an important factor putting rates higher or lower than they otherwise would be whereas the sentence count finds uncertainty referenced in nearly one-third of the sentences in the policy section of the minutes. Incoming data on economic activity were soft, and the Committee was uncertain over the degree to which the economy was weakening. At the same time, they had a good deal of uncertainty on whether their expected decline in inflation – which was running uncomfortably high at the time – actually would materialize. In the end, they only removed the bias in the statement towards further tightening, and did not adjust policy one way or the other in response to the conflicting uncertainties. Hence the judgmental indicator did not code policy being higher of lower than it otherwise would be due to uncertainty.

At other times, the word count was a more simple misread of the Committee's intentions.

For example, in March 2000 the word count identified an insurance coding since it found the word "ensure" in the policy portion of the minutes. However, this turned out not to be associated with the current policy decision, but a comment with regard to the possible need to increase rates in the future to ensure inflation remains contained, and hence was not coded in our judgmental insurance indicator.

We conclude this section by discussing why we did not attempt to use the minutes to measure any variables for risk management *per se*. The minutes often contain discussions of risks to the Committee's dual mandate goals. But when not accompanied by references to uncertainty or insurance, the risk management language may simply describe policy settings that balance conflicting risks to the outlooks of output and inflation relative to their implicit targets. Such policy moves may just be adjusting the expected losses along output and inflation paths in a balanced fashion as is prescribed by the canonical framework for studying optimal policy under discretion. This was the issue we discussed in our earlier narrative of the 1997-1998 period.[38]

Some references to risk do appear to refer to circumstances that moved policy actions away from the norm. Nevertheless we coded our indicator variables to zero for some of these meetings as well. March 2008 provides an example. At that meeting the staff was projecting a mild recession followed by a fair-sized recovery. This forecast was conditioned on a relatively aggressive cut in the FFR; and in the event, the Committee lowered rates be 25 basis points more than the staff had assumed. The minutes state:

> . . . most members judged that a substantial easing in the stance of monetary pol-
> icy was warranted at this meeting. The outlook for economic activity had weak-
> ened considerably since the January meeting, and members viewed the downside
> risks to economic growth as having increased. Indeed, some believed that a

---

[38]Indeed, for much of our sample period, the Committee discussed risks about the future evolution of output or inflation relative to target in order to signal a possible bias in the direction of upcoming rate actions. For example, in the July 1997 meeting described earlier, the minutes indicate members ". . . wanted to retain the existing asymmetry toward restraint . . . An asymmetric directive was consistent with their view that the risks clearly were in the direction of excessive demand pressures . . . " Since the Committee delayed tightening at this meeting, this "risk" reference communicated that the risks to price stability presented by the baseline outlook would likely eventually call for rate increases. It is not a reference that variance or skewness in the distribution of possible inflation outcomes should dictate some non-standard policy response.

prolonged and severe economic downturn could not be ruled out ...

The Committee's actions and this reference to risk appears, like the January 2001 example cited above, to be a concern about nonlinear recessionary dynamics or non-quadratic losses. According to our discussion of theoretical rationales for risk management in Section 3.1, such factors should elicit a non-standard policy response. However, while there were references to elevated uncertainty about growth and inflation in the March 2008 minutes it was not clear to us that policy was tilted one way or another in response to these uncertainties. Therefore we coded our indicator variables to zero for this meeting.

Given these disparate uses of the term "risk" in the minutes, we do not confine our empirical analysis to judgmental coding or sentence counts from the FOMC minutes. We consider other variables not specific to the minutes to try to uncover when variance or skewness about point forecasts or unusual surprises in the outlook generated an atypical policy response. We discuss how we do this below in Section 4.2.

4 Econometric Evidence on Risk Management

We have uncovered clear evidence that risk management considerations have been influential in determining the FOMC's policy stance. This suggests that a proposal to guide policy today by managing the risk of returning to the ZLB is not inconsistent with the words of the Committee. But, it is not clear at this stage whether risk management has had a material impact on the FOMC's FFR choices. Have the words of the FOMC been reflected in its deeds? If the answer to this question is "no" then any proposal to incorporate risk management into policy-making in the current environment would be harder to justify. Therefore in this section we quantify the effect risk management has had on monetary policy in the pre-ZLB era.

We estimate monetary policy reaction functions of the kind studied in Clarida, Gali, and Gertler (2000) and many other papers. These have the FFR set as a linear function of output gap and inflation forecasts; there is no role for risk management unless it feeds directly into

43

these forecasts. To quantify the role of risk management beyond any direct influence it has on forecasts we add to the reaction function variables that proxy for it. Since we focus on a pre-ZLB sample a test for the statistical significance of a given proxy's coefficient amounts to a test of whether the FOMC has followed the normative prescriptions of the theories discussed in Section 3.1.[39] However an insignificant coefficient cannot be interpreted as evidence against a role for risk management precisely because it can influence forecasts.[40] Considering a broad array of proxies we find plenty of evidence that risk management has had a statistically and economically significant impact on monetary policy beyond any direct effects on the forecast.[41]

4.1 Empirical Strategy

Let R_t^* denote the *notional target* for FFR in period t. We assume the FOMC uses the following rule for setting this target:

$$R_t^* = R^* + \beta \left(E_t \left[\pi_{t,k} \right] - \pi^* \right) + \gamma E_t \left[x_{t,q} \right] + \mu s_t, \tag{10}$$

where $\pi_{t,k}$ denotes the average annualized inflation rate from t to $t+k$, π^* is the FOMC's target for inflation, $x_{t,q}$ is the average output gap from t to $t+q$, s_t is a risk management proxy and E_t is the expectations operator conditional on information available to the FOMC at date t. The coefficients β, γ and μ are assumed to be fixed over time. R^* is the central bank's desired nominal rate when inflation is at target, the output gap is closed and uncertainty does not influence policy, $\mu = 0$. Assume the average output and inflation gaps both equal zero. Furthermore suppose the FOMC acts as if the natural rate is constant and out of its

---

[39]A given theory could be correct and our tests indicate no effects of uncertainty on policy because the Fed has not conducted optimal policy. Testing the positive implications of the theories discussed in Section 3.1 is beyond the scope of this paper.

[40]For example, in our discussion of the expectations channel in Section 2.1 uncertainty over the future natural rate has a direct impact on optimal policy, but the expected value of the output gap is a sufficient statistic for this uncertainty.

[41]There is a large literature that examines non-linearities in policy reaction functions (see Gnabo and Moccero (2014), Mumtaz and Surico (2015), and Tenreyro and Thwaites (2015) for reviews of this literature and recent estimates), but surprisingly little work that speaks directly to risk management. We discuss the related literature below.

control. Then $R^* = r^* + \pi^*$, where r^* is the natural rate.[42]

Our estimation equation embodies two additional assumptions. First, the FOMC has a preference for interest rate smoothing and so does not choose the FFR to hit its notional target instantaneously. Second, the FOMC does not have perfect control over interest rates. This motivates the following specification for the *effective target* FFR, R_t

$$R_t = (1 - \bar{a})R_t^* + A(L)R_{t-1} + \upsilon_t \tag{11}$$

where $A(L)$ is a polynomial in the lag operator L given by

$$A(L) = \sum_{j=0}^{N-1} a_{j+1}L^j \ \text{ and } \ 0 \leq \bar{a} \equiv \sum_{j=0}^{N-1} a_{j+1} < 1,$$

with N denoting the number of FFR lags. The variable υ_t is a mean zero and serially independent interest rate shock. Combining (10) and (11) yields our estimation equation:

$$R_t = b_0 + b_1 E_t \left[\pi_{t,k}\right] + b_2 E_t \left[x_{t,q}\right] + A(L)R_{t-1} + b_3 s_t + \upsilon_t. \tag{12}$$

where b_i, $i = 0, 1, 2, 3$ are simple functions of $\bar{a}$, β, γ, μ, r^* and π^*.

We use the publicly available Board staff forecasts of core CPI inflation (in percentage points) and the output gap (percentage point deviations of real GDP from its potential) to measure $E_t \left[\pi_{t,k}\right]$ and $E_t \left[x_{t,q}\right]$ with $k = q = 3$.[43] These forecasts are available for all of the eight FOMC meetings per year so when estimating (12) with our meeting-based proxies we use all these observations.[44] We study our other proxies at the quarterly frequency. In these cases our results are based on staff forecasts corresponding to FOMC meetings closest to the middle of each quarter. We measure R_t in meeting-based cases with the FFR announced

---

[42]There is no presumption that (10) reflects optimal policy and so assuming a constant natural rate is not inconsistent with our theoretical analysis. We explored using forecasted growth in potential output derived from Board staff forecasts to proxy for the natural rate and found this did not affect our results.

[43]Gnabo and Moccero (2014) estimate policy reactions functions using staff forecasts as well. The forecasts are obtained from the Federal Reserve Bank of Philadelphia public web site.

[44]In this estimation the time between meetings is held constant even though this is not true in practise. We account for this discrepancy when we calculate standard errors by allowing for heteroskedasticity.

at the end of the meeting and in the quarterly cases with the average FFR over the 30 days following a meeting. Because our measures of $E_t[\pi_{t,k}]$ and $E_t[x_{t,q}]$ are based solely on information available before an FOMC it follows that we can obtain consistent estimates of β, γ and μ by estimating (12) by ordinary least squares, as long as there are sufficient lags in R_t to ensure that the errors υ_t are serially uncorrelated.[45]

As discussed above we test the null hypothesis that risk management does not impact the setting of the FFR by estimating (12) and determining whether μ is significantly different from zero. We do not allow for the coefficients on the macroeconomic forecasts in the reaction function to depend on uncertainty as suggested by the work of Brainard (1967) and others. However as we (will) show in the appendix that if these coefficients are linear functions of uncertainty then the null hypothesis $\mu = 0$ encompasses the hypothesis that uncertainty does not affect the reaction function coefficients.

4.2 Proxies for Risk Management

In addition to our FOMC-minutes-based variables we consider several proxies for risk management that do not rely on our interpretation of what the words in the FOMC meetings mean. Two of these variables are derived from the Board staff's forecast seen by the FOMC at their regular meetings and so we study them along with our other meeting-based proxies; the remaining ones are studied at the quarterly frequency. These latter variables can be divided into two groups based on our assessment of whether they primarily reflect variance or skewness in the forecast.

The two meeting-based proxies involve revisions to the Board staff's forecasts for the output gap (frGap) and core CPI inflation (frInf). The revisions correspond to changes in the forecast for the same one year period made between meeting t and $t-1$. These variables allow us to assess whether shocks that move forecasts a large amount elicit an extra policy response. If the Committee was only worried about the effects of unusual events on its point

---

[45]We make no attempt to address the impact of the possibility of hitting the ZLB in our estimation. See Chevapatrakul, Kim, and Mizen (2009) and Kiesel and Wolters (2014) for papers that do this.

forecast, then the post-shock projection of the output or inflation gaps would be sufficient to describe the policy setting. But, if a large forecast revision also signals an asymmetric weight on outcomes in the direction of the shock and the FOMC wanted to insure against those outcomes, we could observe effects of revisions on the FFR beyond their influence on the point forecast.

Two of the quarterly proxies are based on financial market data: VIX and Spread. VIX is the well-known measure of market participants' expectations of volatility of the S&P 500 stock index over the next 30 day period.[46] This variable possibly confounds uncertainty due to financial factors that could be unrelated to the outlook for the economy. However, the S&P 500 reflects expectations of earnings so that VIX should measure market participants' uncertainty about the outlook for the economy over horizons relevant to the FOMC.[47] Spread is simply the difference between the quarterly average of daily yields on BAA corporate bonds and 10 year Treasury bonds. Gilchrist and Zakrajšek (2012) demonstrate that this variable measures information on private sector default risk plus other factors that may indicate downside risks to economic growth.[48]

The remaining proxies we look at are based on the Survey of Professional Forecasters (SPF). The SPF surveys professional forecasters about their point forecasts of GDP growth and GDP deflator inflation and their probability distributions for these forecasts.[49] We use both kinds of information to construct measures of variance and skewness in the economic outlook one year ahead. We measure variance using the median among forecasters of the standard deviations calculated from each individual probability distribution (vGDP and

---

[46]This is the Chicago Board Options Exchange Market Volatility Index obtained from the Haver Analytics database.

[47]Bekaert, Hoerova, and Lo Duca (2013) study the impact of innovations to VIX on the FFR in a VAR setting. They find a positive innovation to VIX leads to looser policy although their findings are not very robust and only weakly significant. Gnabo and Moccero (2014) find that policy responds more aggressively and the degree of inertia in policy is lower in periods of high economic risk as measured by VIX.

[48]Alcidi, Flamini, and Fracasso (2011), Castelnuovo (2003) and Gerlach-Kristen (2004) consider the role of Spread in estimated monetary policy reaction functions.

[49]The forecast distributions are for growth and inflation in the current and following year. We use D'Amico and Orphanides (2014)'s procedure to translate these distributions into distributions for GDP growth and inflation over the next four quarters. Note that the bins forecasters are asked to put probability mass on are 1 percentage point wide, so moments calculated using them may contain substantial measurement error.

vInf) and the interquartile range of point forecasts across individuals (PFvGDP and PFv-Inf.)[50] PFvGDP and PFvInf are properly thought of as measuring forecaster disagreement, but there is a large literature that uses forecaster disagreement as a proxy for variance.[51] To measure skewness we use the median of the individual forecasters' mean less the mode (sGDP and sInf) and the difference between the mean and the mode of the cross-forecaster distribution of point forecasts (PFsGDP and PFsInf).

All samples run to the end of 2008 due to the onset of the ZLB, but begin at different dates according to the number of observations available for each variable. The benchmark start date is determined by the beginning of Alan Greenspan's tenure as Chairman of the FOMC in 1987, but later dates are used for two sets of variables due to limitations of the data. The first observation for the meeting-based indicators is the first FOMC meeting of 1993 for the reason discussed in Section 3. When considering the proxies based on forecast distributions of individual forecasters from the SPF the first observation is 1992q1 due to a discrete change in SPF methodology that occurs then.[52]

Tables 4 and 5 display various summary statistics for Board staff forecasts of inflation and the output gap and the various proxies for risk management at the meeting and quarterly frequencies. Several entries in these tables are worth highlighting. First, the staff's forecasts of the output gap and inflation are essentially uncorrelated. Second, Spread is strongly negatively correlated with the forecast variables, consistent with it indicating downside risks to the economy. Third, the remaining variables mostly display weak correlations with the staff forecasts, with two notable exceptions. Inflation forecast disagreement (PFuInf) has a large negative correlation with the output gap forecast: periods when the staff's outlook for activity is deteriorating often correspond to periods when there is a large amount of

---

[50]Gnabo and Moccero (2014) study the effects of PFvInf on monetary policy but do not find statistically significant effects.

[51]However there is not a consensus about how good a proxy it is. See Baker, Bloom, and Davis (2015) for a recent review of the relevant literature. Note that we do not use the measure of uncertainty constructed by Baker et al. (2015) in our analysis since it includes information reflecting uncertainty about monetary policy.

[52]In 1992 the SPF narrows the bins it uses to summarize the forecast probability distributions of individual forecasters. See D'Amico and Orphanides (2014) and Andrade, Ghysels, and Idier (2013) for attempts to address this change in bin sizes.

Table 4: Meeting-based variables' summary statistics

Variable	Obs.	Mean	Std. Dev.	Min	Max	Correlation with forecast of	
						Inflation	Output gap
Inflation forecast	167	2.9	0.98	1.3	5.60	1.00	
Output gap forecast	167	-0.42	1.74	-4.97	3.08	-0.02	1.00
hUnc	128	-0.13	0.48	-1	1	-0.23	-0.33
hIns	128	-0.06	0.33	-1	1	0.18	0.15
mUnc	128	2.92	4.8	0	30.77	-0.06	0.14
mIns	128	0.83	2.45	0	16.67	-0.10	0.08
frInf	167	-0.02	0.21	-0.68	0.70	0.08	0.04
frGap	167	-0.02	0.42	-2	0.83	0.04	0.29

Note: See the text for a description of the sample periods.

Table 5: Quarterly variables' summary statistics

Variable	Obs.	Mean	Std. Dev.	Min	Max	Correlation with forecasts of	
						Inflation	Output Gap
Inflation forecast	83	2.9	0.98	1.33	5.32	1.00	
Output gap forecast	83	-0.43	1.72	-4.4	3.08	-0.02	1.00
VIX	83	20.48	7.92	10.58	62.09	-0.15	0.06
vInf	68	0.74	0.06	0.6	0.9	-0.22	-0.08
vGDP	68	0.9	0.12	0.67	1.3	-0.22	0.22
PFvInf	83	0.6	0.18	0.24	1.1	0.26	-0.36
PFvGDP	83	0.71	0.25	0.3	1.62	0.30	-0.04
Spread	83	2.11	0.66	1.37	5.60	-0.37	-0.34
sInf	68	0.05	0.08	-0.12	0.3	0.23	-0.12
sGDP	68	-0.1	0.19	-0.54	0.47	-0.10	-0.48
PFsInf	83	0.06	0.21	-0.5	0.51	0.01	-0.23
PFsGDP	83	0.29	0.27	-0.5	0.9	-0.31	0.23

Note: See the text for a description of the sample periods.

disagreement about the outlook for inflation. In addition, skewness in forecasters' GDP forecasts (sGDP) is strongly *negatively* correlated with the outlook for activity. This suggests private sector forecasters are slower than the Board staff to react to downside risks to activity.

Table 6: Cross-correlations of FOMC-meeting-based risk management variables

Variables	hUnc	hIns	mUnc	mIns	frInf
hIns	-0.05				
mUnc	-0.07	0.02			
mIns	-0.13	-0.09	0.04		
frInf	0.10	0.05	-0.07	0.06	
frGap	-0.11	0.08	0.05	0.11	0.22

Table 7: Cross-correlations of quarterly variance and skewness variables

Variables	VIX	uInf	uGDP	PFuInf	PFuGDP	Spread	sInf	sGDP	PFsInf
vInf	0.04								
vGDP	0.40	0.40							
PFvInf	0.15	0.29	0.03						
PFvGDP	0.44	0.16	0.37	0.34					
Spread	0.79	0.15	0.32	0.25	0.37				
sInf	-0.27	0.29	-0.16	0.08	-0.14	-0.18			
sGDP	0.21	-0.09	-0.04	0.25	0.16	0.43	-0.15		
PFsInf	-0.29	0.10	-0.24	0.16	-0.13	-0.08	0.04	0.15	
PFsGDP	-0.01	0.04	0.10	-0.22	0.00	0.01	-0.08	-0.17	-0.19

Tables 6 and 7 display cross-correlations of the FOMC-meeting-based and quarterly proxies, respectively. As suggested by Figures 4 and 5 the human and machine coded indicator variables for inflation and uncertainty are essentially uncorrelated. These variables also appear unrelated to the forecast revision variables. Several correlations among the quarterly proxies are worth noting. Forecaster uncertainty and disagreement about the GDP growth outlook (vGDP and PFvGDP) are both positively correlated with VIX and Spread, suggesting the financial variables are good indicators of uncertainty about the activity outlook. Interestingly the correlation of VIX with the inflation asymmetry variables (sInf and PFsInf) are both negative: when markets perceive a lot of uncertainty in the stock market the inflation outlook is skewed to the downside. The correlation of vGDP with vInf and PFvGDP

with PFvInf are both fairly large suggesting uncertainty about inflation and GDP often move together. Finally, the correlations of the corresponding forecaster uncertainty and disagreement variables (vGDP with PFvGDP and vInf with PFvInf) are somewhat large too. Evidently disagreement among forecasters is similar to the median amount of uncertainty they see.

4.3 Policy Rule Findings

Table 8 shows our policy rule estimates with and without the various FOMC-meeting-based variables; Tables 9 and 10 show estimates with and without the variance and skewness variables. The tables have the same layout with the first columns showing the policy rule estimates without any risk management variables and the other columns show the results of estimating the policy rule adding one of the risk management variables at a time with the indicated coefficient estimate corresponding to μ in (10). In the policy rules without any risk management variables the coefficient on inflation (β) is about 1.8 and on the output gap (γ) is about 0.8. These estimates are highly significant and are similar to estimated forecast-based policy rules in the literature. Introducing one of the proxies for risk management typically moves the reaction coefficients very little with a couple of exceptions noted below.

Table 8 indicates that the human coding of uncertainty (hUnc) is statistically significant at the 5% level. The coefficient indicates that on average when uncertainty has shaded the policy decision above or below what would be dictated by the forecast (as determined by our analysis of the FOMC minutes) it has moved the notional target (see 10) by 40 basis points (bps). With interest rate smoothing the immediate impact is much smaller; the 95% confidence interval is 2-14 bps. The machine coding of uncertainty is significant at the 10% level but the effect is small – a one standard deviation increase in the number of sentences we associate with uncertainty raises the notional target by 14bps.[53] The insurance indicators (hIns and mIns) are not significant, but the point estimate of the human coded variable

---

[53] The machine coded variables are harder to interpret than the human coded indicators since they do not account for whether the mentions of uncertainty or insurance shaded policy up or down. It would be interesting to examine whether a more sophisticated parsing of the words would yield stronger results.

indicates its affects are similar to its uncertainty counterpart. Perhaps the most striking result in Table 8 is the large and highly significant coefficient on the output gap forecast revision variable (frGap).[54] The point estimate indicates a one standard deviation (42 bps) positive surprise in the forecast raises the notional target by 63 bps over and above the impact this surprise has on the forecast itself; the confidence interval is 18-108 bps.

Table 8: FOMC meeting variables in monetary policy rules

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
β	1.77***	1.95***	1.86***	1.90***	1.89***	1.76***	1.77***
	(.11)	(.16)	(.17)	(.17)	(.17)	(.13)	(.11)
γ	.79***	.88***	.85***	.83***	.85***	.72***	.79***
	(.06)	(.05)	(.05)	(.05)	(.05)	(.07)	(.06)
hUnc		.40**					
		(.16)					
hIns			.48				
			(.45)				
mUnc				.03*			
				(.01)			
mIns					-.0003		
					(.03)		
frGap						1.50***	
						(.54)	
frInf							.32
							(.66)
Obs.	167	128	128	128	128	167	167
LM	.02	.07	.59	.58	.31	.34	.06

Note: The superscripts ***, ** and * indicate statistical significance at the 1, 5 and 10 percent levels respectively. Standard errors are robust to heteroskedasticity. All models include 5 lags of the FFR. The entries in the row labelled "LM" are p-values of Durbin's test for the null hypothesis of no serial correlation in the regression residuals up to fifth order.

Table 9 shows strong evidence that risk management has shaded policy away from that predicted by forecasts alone. The coefficient on VIX is highly significant and enters with a negative sign. A one standard deviation increase in VIX lowers the notional target FFR

---

[54]The magnitude and significance of this coefficient is partly driven by the observations from 2008, but not entirely. Ending the sample in 2007 lowers the point estimate to 0.9 but this remains significant at the 5% level.

Table 9: Variance in monetary policy rules

	(1)	(2)	(3)	(4)	(5)	(6)
β	1.79***	1.74***	2.21***	2.13***	1.82***	1.91***
	(.12)	(.12)	(.17)	(.16)	(.12)	(.14)
γ	.79***	.84***	.78***	.77***	.76***	.80***
	(.06)	(.06)	(.07)	(.06)	(.07)	(.06)
VIX		-.05***				
		(.01)				
vInf			3.47**			
			(1.67)			
vGDP				.26		
				(.98)		
PFvInf					-.54	
					(.67)	
PFvGDP						-1.43***
						(.53)
Obs.	83	83	68	68	83	83
LM	.70	.80	.71	.59	.70	.92

Note: Superscripts ***, ** and * indicate statistical significance at the 1, 5 and 10 percent levels respectively. Standard errors are robust to heteroskedasticity. All models include 2 lags of the FFR. Entries in the row labelled "LM" are p-values of Durbin's test for the null hypothesis of no serial correlation in the regression residuals up to second order.

by 40bps. Consistent with this, the coefficient on GDP forecast uncertainty as measured using point forecast dispersion (PFvGDP) is also highly significant and indicates roughly the same affect on the target. The coefficient on the variable measuring how forecasters view the uncertainty in their inflation forecasts (vInf) is significant. In this case uncertainty shades the policy higher; a standard deviation increase in vInf raises the notional FFR target by about 25bps.

Similarly strong evidence that skewness matters for policy decisions is indicated in Table 10. The coefficients on Spread (the interest rate spread indicator of downside risks to activity), sInf (skewness in the outlook for inflation measured from forecasters' own forecast distributions) and PFsInf (skewness in the inflation outlook measured across point forecasts) all enter into the policy rule significantly. An increase in downside risks to activity lowers while an increase in upside risks to inflation raises the FFR. The effects seem large. Standard

Table 10: Skewness in monetary policy rules

	(1)	(2)	(3)	(4)	(5)	(6)
β	1.79***	1.59***	2.02***	2.09***	1.80***	1.75***
	(.12)	(.12)	(.16)	(.16)	(.11)	(.13)
γ	.79***	.71***	.80***	.74***	.87***	.80***
	(.06)	(.06)	(.07)	(.08)	(.08)	(.07)
Spread		-.82***				
		(.21)				
sInf			2.80**			
			(1.18)			
sGDP				-.79		
				(.58)		
PFsInf					1.80***	
					(.64)	
PFsGDP						-.46
						(.47)
Obs.	83	83	68	68	83	83
LM	.70	.94	.34	.67	.84	.72

Note: Superscripts ***, ** and * indicate statistical significance at the 1, 5 and 10 percent levels respectively. Standard errors are robust to heteroskedasticity. All models include 2 lags of the FFR. Entries in the row labelled "LM" are p-values of Durbin's test for the null hypothesis of no serial correlation in the regression residuals up to second order.

deviation increases in these proxies change the notional target by -54, 22 and 38 bps, respectively. The point estimates for skewness in the GDP outlook come in with unexpectedly negative signs (perceived upside risks to the growth suggest lowering the FFR). However these coefficients are relatively small and insignificant (standard deviation changes translate to no more than 15 bp changes in the notional target.) The standard errors are large enough that the expected sign cannot be ruled out.

Taken together, these results suggest that risk management concerns broadly conceived have had a statistically and economically significant impact on policy decisions over and above how those concerns are reflected in point forecasts. Risk management does not just appear in the words of the FOMC it is apparently reflected in their deeds as well.

5 Conclusion

Our analysis has so far ignored two reputational issues that may be relevant to the lift off calculus. A policymaker must take into account the effect that shocks might have on her reputation; in particular, policymakers may face large costs of reversing a decision. Empirically, it is well known that central banks tend to go through "tightening" and "easing" cycles which in turn induce substantial persistence in the short-term interest rate. One reason why policymakers might be reluctant to reverse course is that it would damage their reputation, perhaps because the public would revise its confidence in the central bank's ability to understand and stabilize the economy. With high uncertainty, this reputation element would lead to more caution. In the case of liftoff, it would argues for a longer delay in raising rates to avoid the reputational costs of a reversion back to the ZLB.

Another reputational concern is the signal the public might infer about the central bank's commitments to its stated policy goals if liftoff occurred with output or inflation still far below target. Large gaps on their own pose no threat to credibility if the public is confident that the economy is on a path to achieve the policy targets in a reasonable period of time and that the central bank is willing to accommodate this path. However, if there is elevated uncertainty over the strength of the economy, or a view that risks were skewed to the downside, then the public might believe that a risk management approach should delay liftoff. And as we found, the larger the current gaps, the more relevant these risk management concerns. Accordingly, early liftoff might be construed as backing away from appropriate risk management, and hence a less-than-enthusiastic endorsement of the ultimate policy targets by the central bank. This could work in the opposite direction for an economy in which the uncertainties or asymmetries underlying the risk management approach dictated an aggressive tightening move to guard against inflation. Either way, the central bank's reputation as having a credible commitment to achieving its policy goals could brought into question.

References

Adam, K. and R. M. Billi (2007). Discretionary monetary policy and the zero lower bound on nominal interest rates. *Journal of Monetary Economics 54*(3), 728–752.

Alcidi, C., A. Flamini, and A. Fracasso (2011). Policy regime changes, judgment and taylor rules in the greenspan era. *Economica 78*, 89–107.

Andrade, P., E. Ghysels, and J. Idier (2013). Tails of inflation forecasts and tales of monetary policy. UNC Kenan-Flagler Research Paper No. 2013-17.

Baker, S. R., N. Bloom, and S. J. Davis (2015). Measuring economic policy uncertainty. Stanford University manuscript.

Barlevy, G. (2011). Robustness and macroeconomic policy. *Annual Review of Economics 3*, 1–24.

Barsky, R., A. Justiniano, and L. Melosi (2014). The natural rate and its usefulness for monetary policy making. *American Economic Review 104*(4), 37–43.

Basu, S. and B. Bundick (2013). Downside risk at the zero lower bound. Boston College Manuscript.

Bekaert, G., M. Hoerova, and M. Lo Duca (2013). Risk, uncertainty and monetary policy. *Journal of Monetary Economics 60*, 771–788.

Bernanke, B. S. (2005). The changing policy landscape. In *Monetary Policy Since the Onset of the Crisis*, Economic Policy Symposium, pp. 1–22. Federal Reserve Bank of Kansas City.

Blinder, A. (1998). *Central Banking in Theory and Practice*. Cambridge, MA: MIT Press.

Bloom, N. (2009). The effect of uncertainty shocks. *Econometrica 77*(3), 623–685.

Bomfin, A. and L. Meyer (2010). Quantifying the effects of fed asset purchases on treasury yields. Macroeconomics Advisors, Monetary Policy Insights: Fixed Income Focus.

Born, B. and J. Pfeifer (2014). Policy risk and the business cycle. *Journal of Monetary Economics 68*, 68–85.

Brainard, W. (1967). Uncertainty and the effectiveness of policy. *American Economic Review 57*(2), 411–425.

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 Spring*, 1–54.

Castelnuovo, E. (2003). Taylor rules, omitted variables, and interest rate smoothing in the US. *Economics Letters 81*, 55–59.

Chen, H., V. Curida, and A. Ferrero (2012). The macroeconomic effects of large scale asset purchase programmes. *Economic Journal 122*, 289–315.

Chevapatrakul, T., T. Kim, and P. Mizen (2009). The Taylor principle and monetary policy approaching a zero bound on nominal rates: Quantile regression results for the united states and japan. *Journal of Money, Credit and Banking 41*(8), 1705–1723.

Christiano, L., M. Eichenbaum, and C. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45.

Clarida, R., J. Gali, and M. Gertler (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics CXV*(1), 147–180.

Coibion, O., Y. Gorodnichenko, and J. Wieland (2012). The optimal inflation rate in new keynesian models: Should central banks raise their inflation targets in light of the zero lower bound? *Review of Economic Studies 79*(4), 1371–1406.

D'Amico, S. and T. King (2013). Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. *Journal of Financial Economics 108*, 425–448.

D'Amico, S. and T. King (2015). Policy expectations, term premia, and macroeconomic performance. Federal Reserve Bank of Chicago manuscript.

D'Amico, S. and A. Orphanides (2014). Inflation uncertainty and disagreement in bond risk premia. Chicago Fed working paper 2014-24.

Dolado, J. J., P. R. María-Dolores, and M. Naveira (2005). Are monetary policy reaction functions asymmetric?: The role of nonlinearity in the phillips curve. *European Economic Review 49*(2), 485–503.

Dolado, J. J., P. R. María-Dolores, and F. J. Ruge-Murcia (2004). Nonlinear monetary policy rules: Some new evidence for the us. *Studies in Nonlinear Dynamics and Econometrics 8*(3).

Egertsson, G. B. and M. Woodford (2003). The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity 2003*(1), 139–211.

Eggertsson, G. B. and B. Pugsley (2009). The mistake of 1937: A general equilibrium analysis. *Monetary and Economic Studies 24*(S-1), 16–37.

Eichenbaum, M. and J. D. Fisher (2007). Esimating the frequency of price re-optimization in calvo-style models. *Journal of Monetary Economics 54*(7), 2032–2047.

Engen, E. M., T. T. Laubach, and D. Reifschneider (2015). The macroeconomic effects of the federal reserve's unconventional monetary policy. Finance and Ecnomics Discussion Series 2012-5, Board of Governors of the Federal Reserve System.

English, W. B., D. López-Salido, and R. Tetlow (2013). The federal reserve's framework for monetary policy – rcent changes and new questions. Federal Reserve Board, Finance and Economics Discussion Series 2013-76.

Evans, C. L. (2014). Patience is a virtue when normalizing monetary policy. Text of speech to Peterson Institute for International Economics.

Fernándex-Villaverde, J., P. Guerró-Quintana, K. Kuester, and J. Rubio-Ramírez (2012). Fiscal volatility shocks and economic activity. Philladelphia Fed Working Paper No. 11-32/R.

Fuhrer, J. C. (2000). Habit formation in consumption and its implications for monetary-policy models. *America Economic Review 90*(3), 367–390.

Gagnon, J., M. Raskin, J. Remache, and B. Sack (2010). Optimal fiscal and monetary policy with occasionally binding zero bound constraints. Federal Reserve Bank of New York Staff Report No. 441.

Gali, J. (2008). *An Introduction to the New Keynesian Framework*. Princeton, NJ: Princeton University Press.

Gali, J. and M. Gertler (1999). Inflation dynamics: a structural econometric analysis. *Journal of Monetary Economics 44*, 192–222.

Gerlach-Kristen, P. (2004). Interest rate smoothing: Monetary policy inertia or unobserved variables. *Contributions in Macroeconomics 4*(1), 1534–6005.

Gilchrist, S. and E. Zakrajšek (2012). Credit spreads and business cycle fluctuations. *American Economic Review 102*(4), 1692–1720.

Gnabo, J. and D. N. Moccero (2014). Risk management, nonlinearity and agressiveness in monetary policy: The case of the US Fed. *Journal of Banking & Finance*.

Greenspan, A. (2004). Risk and uncertainty in monetary policy. *American Economic Review 94*(2), 33–40.

Hamilton, J. and J. Wu (2010). The effectiveness of alternative monetary policy tools in a zero lower bound environment. UCSD manuscript.

Hamilton, J. D., E. S. Harris, J. Hatzius, and K. D. West (2015). The equilibrium real funds rate: Past, present and future. UCSD manuscript.

Hansen, L. and T. Sargent (2008). *Robustness*. Princeton, NJ: Princeton University Press.

Hansen, S. and M. McMahon (2014). First impressions matter: Signaling as a source of policy dynamics. University of Warwick manuscript.

Johannsen, B. K. (2014). When are the effects of fiscal uncertainty large? Finance and Ecnomics Discussion Series 2014-40, Board of Governors of the Federal Reserve System.

Kiesel, K. and M. H. Wolters (2014). Estimating monetary policy rules when the zero lower bound on nominal interest rates is approached. Kiel Working Paper No. 1898.

Kiley, M. T. (2012). The aggregate demand effects of short- and long-term interest rates. Finance and Ecnomics Discussion Series 2012-54, Board of Governors of the Federal Reserve System.

Kilian, L. and S. Manganelli (2008). The central banker as a risk manager: Estimating the federal reserve's preferences under greenspan. *Journal of Money, Credit and Banking 40*(6).

Krishnamurthy, A. and A. Vissing-Jorgensen (2013). The ins and outs of lsaps. Kansas City Federal Reserve Symposium on Global Dimensions of Unconventional Monetary Policy.

Krugman, P. R. (1998). It's baaack: Japan's slump and the return of the liquidity trap. *Brookings Papers on Economic Activity Fall*, 137–187.

Laubach, T. and J. C. Williams (2003). Measuring the natural rate of interest. *The Review of Economics and Statistics 85*(4), 1063–1070.

Laxton, D., D. Rose, and D. Tambakis (1999). The U.S. phillips curve: The case for asymmetry. *Journal of Economic Dynamics & Control 23*, 1459–1485.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995). *Microeconomic Theory*. Oxford, United Kingdom: Oxford University Press.

Mumtaz, H. and P. Surico (2015). The transmission mechanism in good and bad times. Forthcoming *International Economic Review*.

Nakata, T. (2013a). Optimal fiscal and monetary policy with occasionally binding zero bound constraints. Finance and Ecnomics Discussion Series 2013-40, Board of Governors of the Federal Reserve System.

Nakata, T. (2013b). Uncertainty at the zero lower bound. Finance and Ecnomics Discussion Series 2013-09, Board of Governors of the Federal Reserve System.

Nakata, T. and S. Schmidt (2014). Conservatism and liquidity traps. Finance and Ecnomics Discussion Series 2014-105, Board of Governors of the Federal Reserve System.

Nakov, A. A. (2008). Optimal and simple monetary policy rules with zero floor on the nominal interest rate. *International Journal of Central Banking 4*(2), 73–127.

Orphanides, A. and J. C. Williams (2002). Robust monetary policy rules with unknown natural rates. *Brookings Papers on Economic Activity Fall*, 63–145.

Reifschneider, D. and J. C. Williams (2000, November). Three lessons for monetary policy in a low-inflation era. *Journal of Money, Credit, and Banking 32*(4), 936–966.

Romer, C. D. and D. H. Romer (1989). Does monetary policy matter? a new test in the spirit of friedman and schwartz. In O. Blanchard and S. Fischer (Eds.), *NBER Macroeconomics Annual 2005*, Volume 4.

Rudebusch, G. and L. E. Svensson (1999). Policy rules for inflation targeting. In J. B. Taylor (Ed.), *Monetary Policy Rules*. Chicago, IL: University of Chicago Press.

Rudebusch, G. D. (2001). Is the fed too timid? monetary policy in an uncertain world. *The Review of Economics and Statistics 83*(2), 203–217.

Sack, B. (2000). Does the fed act gradually? a var analysis. *Journal of Monetary Economics 46*(1), 229–256.

Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *The American Economic Review 97*(3), 586–606.

Surico, P. (2007). The fed's monetary policy rule and U.S. inflation: The case of asymmetric preferences. *The Journal of Economic Dynamics & Control 31*, 305–324.

Svensson, L. and M. Woodford (2002). Optimal policy with partial information in a forward-looking model: Certianty-equivalence redux. Manuscript.

Svensson, L. and M. Woodford (2003). Indicator variables for optimal policy. *Journal of Monetary Economics 50*, 691–720.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy 39*, 195–214.

Tenreyro, S. and G. Thwaites (2015). Pushing on a string: US monetary policy is less powerful in recessions. Manuscript.

Werning, I. (2012). Managing a liquidity trap: Monetary and fiscal policy. MIT Manuscript.

Williams, J. (2013). A defense of moderation in monetary policy. *Macroeconomic Dynamics 38*, 137–150.

Woodford, M. (2003). *Interest and Prices*. Princeton, NJ: Princeton University Press.

Woodford, M. (2012). The changing policy landscape. In *Methods of Policy Accommodation at the Interest-Rate Lower Bound*, Economic Symposium, pp. 185–288. Federal Reserve Bank of Kansas City.

Appendix

A   Optimal policy in the forward-looking model with uncertainty about cost-push shocks

Our previous analysis assumed that the unknown shock that might trigger a binding ZLB at time 1 is the natural real rate. We now consider the case where it is the cost-push inflation shock u_1 : *i.e.* $\rho_t^n = \rho$ for $t \geq 1$, and $u_t = 0$ for all $t \geq 2$, but u_1 is distributed according to the probability density function $f_u(.)$. We assume $E(u_1) = 0$.

To find optimal policy, we again solve the model backward. As before, optimal policy after time 2 is simply $x_t = \pi_t = 0$, which is obtained by setting $i_t = \rho > 0$. At time 1, the ZLB may bind if the cost-push shock is negative enough. Specifically, after seeing u_1, we solve

$$\min_{x_1} \frac{1}{2}\left(\pi_1^2 + \lambda x_1^2\right),$$

$$s.t. \quad :$$

$$\pi_1 = \kappa x_1 + u_1,$$

$$x_1 \leq \frac{\rho}{\sigma}.$$

with the following solution:

- If $u_1 \geq u_1^* = -\frac{\rho}{\sigma}\frac{\lambda+\kappa^2}{\kappa}$, the ZLB does not bind, and optimal policy strikes a balance between the inflation and output gap objectives, as in section 2.1:

$$x_1 = -\frac{\kappa u_1}{\lambda + \kappa^2},$$

$$\pi_1 = \frac{\lambda u_1}{\lambda + \kappa^2}.$$

- If $u_1 < u_1^*$, the ZLB binds, so even though the central bank would like to cut rates more to create a larger boom and hence more inflation, this is not feasible. Mathematically,

$$x_1 = \frac{\rho}{\sigma},$$

$$\pi_1 = \kappa\frac{\rho}{\sigma} + u_1.$$

To calculate optimal policy at time 0, we require expected inflation and output. These are given by

$$E\pi_1 = \int_{-\infty}^{u_1^*} \left(\kappa\frac{\rho}{\sigma} + u\right) f_u(u)du + \frac{\lambda}{\lambda + \kappa^2}\int_{u_1^*}^{\infty} u f_u(u)du,$$

$$= \kappa \frac{\rho}{\sigma} P + \frac{\kappa^2}{\lambda + \kappa^2} M,$$

where $P = \int_{-\infty}^{u_1^*} f_u(u) du$ is the probability that the ZLB binds and $M = \int_{-\infty}^{u_1^*} u f_u(u) du$. Note $M < 0$ since $Eu_1 = 0$. Expected output is similarly

$$Ex_1 = \frac{E\pi_1}{\kappa} = \frac{\rho}{\sigma} P + \frac{\kappa}{\lambda + \kappa^2} M.$$

If there was no ZLB, we would have $E\pi_1 = Ex_1 = 0$. With the ZLB, we do worse on output and inflation when there is a negative enough cost-push shock, and hence $Ex_1 < 0$ and $E\pi_1 < 0$.

This implies that optimal policy at time 0 is affected exactly as in the case of a natural rate uncertainty: (i) the lower expected output gap at time 1 leads to a lower output gap at time 0 through the IS equation; (ii) the lower expected inflation $E\pi_1$ leads to lower output gap at time 0 through higher real rates; (iii) the lower expected inflation finally reduces inflation today. All these lead to looser policy. Formally, the optimal policy problem at time 0 is, given shocks ρ_0^n, u_0, to solve

$$\min_{x_0} \frac{1}{2} \left(\pi_0^2 + \lambda x_0^2 \right),$$

$$s.t. \quad : \quad x_0 \leq \frac{\rho_0^n}{\sigma} + Ex_1 + \frac{E\pi_1}{\sigma},$$

$$\pi_0 = \beta E\pi_1 + \kappa x_0 + u_0.$$

The solution is the following. Define

$$\rho_0^* = -\sigma \left(\frac{\rho}{\sigma} P + \frac{\kappa}{\lambda + \kappa^2} M \right) \left(1 + \frac{\beta \kappa^2}{\lambda + \kappa^2} \right) - \frac{\sigma \kappa}{\lambda + \kappa^2} u_0.$$

If $\rho_0^n \geq \rho_0^*$, then optimal policy is described by

$$x_0 = -\frac{\kappa}{\lambda + \kappa^2} \left(\beta E\pi_1 + u_0 \right),$$

$$\pi_0 = \frac{\lambda}{\lambda + \kappa^2} \left(\beta E\pi_1 + u_0 \right),$$

where $E\pi_1 = \kappa \frac{\rho}{\sigma} P + \frac{\kappa^2}{\lambda + \kappa^2} M$. The appropriate interest rate is

$$i_0 = \sigma \left(\frac{\kappa}{\lambda + \kappa^2} \beta E\pi_1 + Ex_1 + u_0 \right) + E\pi_1 + \rho_0^n,$$

so that lower $E\pi_1$ and lower Ex_1 require lower i_0.

If $\rho_0^n < \rho_0^*$, then $i_0 = 0$, and $x_0 = \frac{\rho_0^n}{\sigma} + Ex_1$, and $\pi_0 = (1 + \beta)\kappa Ex_1 + \kappa \frac{\rho_0^n}{\sigma} + u_0$. We can summarize the results in the following proposition:

Proposition 3 *Suppose the uncertainty is about cost-push shocks. Then: (1) optimal policy is looser today when the probability of a binding ZLB tomorrow is positive; (2) optimal policy is independent of the distribution of the cost-push shock tomorrow u_1^n over values for which the ZLB does not bind, i.e. of $\{f_u(u)\}_{u \geq u^*}$; only $\{f_u(u)\}_{u < u^*}$ is relevant, and only through the sufficient statistics $\int_{-\infty}^{u^*} f_u(u)du$ and $\int_{-\infty}^{u^*} u f_u(u)du$.*

Because Ex_1 and $E\pi_1$ now depend on $P = \Pr(u \leq u^*)$, one cannot state a general result about mean-preserving spreads, since this probability might fall with uncertainty for some "unusual" distributions. However, if u is normally distributed with mean 0, and given that $u^* < 0$, the result that more uncertainty leads to lower rates today still hold.

An important implication is that the risk that inflation picks up does not affect policy today. If a high u is realized tomorrow, it will be bad; however, there is nothing that policy *today* can do about it. We finally present an example to illustrate our results.

Example 1 *Suppose that u can take two values, $u = +\Delta$ (with probability $1/2$) and $u = -\Delta$ (with probability $1/2$). If Δ is small, then $P = M = 0$, and hence $E\pi_1 = Ex_1 = 0$, and optimal policy is decided taking into account ρ_0^n and u_0 only. If Δ is large enough, then $P = 1/2$, $M = -\Delta/2$, and $Ex_1 = \frac{\rho}{\sigma}\frac{1}{2} - \frac{\kappa}{\lambda + \kappa^2}\Delta$ (which is negative since $-\Delta < u_1^* = \frac{\rho}{\sigma}\frac{\lambda + \kappa^2}{\kappa}$), and $E\pi_1 = \kappa Ex_1$. A higher Δ then reduces $Ex_1, E\pi_1$ and i_0.*

B Calculation of W in the purely backward-looking model

The value function for $t \geq 2$ solves the following Bellman equation, corresponding to a deterministic optimal control problem:

$$V(\pi_{-1}, x_{-1}) = \min_{x,\pi} \frac{1}{2}\left(\pi^2 + \lambda x^2\right) + \beta V(\pi, x),$$

$$s.t. \quad :$$

$$\pi = \xi\pi_{-1} + \kappa x,$$

$$x = \delta x_{-1} - \frac{1}{\sigma}(i - \rho - \pi_{-1}).$$

We use a guess-and-verify method to show that the value function takes the form

$$V(\pi_{-1}, x_{-1}) = \frac{W}{2}\pi_{-1}^2,$$

and that the policy rules are linear: $\pi = g\pi_{-1}$ and $x = h\pi_{-1}$ for two numbers g and h. To verify the guess, solve

$$\min_x \frac{1}{2}\left(1 + \beta W\right)\left(\xi\pi_{-1} + \kappa x\right)^2 + \frac{1}{2}\lambda x^2$$

The first order condition yields

$$(1 + \beta W)\,(\xi \pi_{-1} + \kappa x)\kappa + \lambda x = 0$$

$$x = -\frac{(1 + \beta W)\kappa \xi}{(1 + \beta W)\kappa^2 + \lambda}\pi_{-1},$$

leading to

$$\pi = \frac{\lambda \xi}{(1 + \beta W)\kappa^2 + \lambda}\pi_{-1},$$

which verifies our guess of linear rules. To find W, plug this back in the minimization problem; we look for W to satisfy, for all π_{-1}, :

$$\frac{W}{2}\pi_{-1}^2 = \frac{1}{2}(1 + \beta W)\left(\frac{\lambda}{(1 + \beta W)\kappa^2 + \lambda}\right)^2 \xi^2 \pi_{-1}^2 + \frac{1}{2}\lambda\left(\frac{(1 + \beta W)\kappa}{(1 + \beta W)\kappa^2 + \lambda}\right)^2 \xi^2 \pi_{-1}^2$$

which can be simplified to a simple quadratic equation:

$$\beta \kappa^2 W + W\left(\kappa^2 + \lambda - \beta \lambda \xi^2\right) = \xi^2 \lambda.$$

It is immediate to verify that, if $\lambda > 0$ and $\xi \neq 0$, there are two real roots to this equation, one negative and one positive. The positive root is our solution and is given by the formula:

$$W = \frac{-\left(\kappa^2 + \lambda(1 - \beta \xi^2)\right) + \sqrt{\left(\kappa^2 + \lambda(1 - \beta \xi^2)\right)^2 + 4\lambda \beta \kappa^2}}{2\beta \kappa^2},$$

and we can calculate g and h given W and the formula above for x and π.

C   Proof of Proposition 2

We start with a simple more general result, then we show how our model fits as a special case of this result.

Lemma 1 *Consider the problem*

$$V(\theta) = \max_{x_0} E_\varepsilon J(x_0, \varepsilon),$$

where θ indexes the distribution of ε, and the function J is defined as

$$J(x_0, \varepsilon) = \max_{x_1} F(x_1, x_0, \varepsilon),$$
$$s.t. \quad : \quad x_1 \leq f(x_0) + \varepsilon,$$

where F is quadratic (with $F_{11} < 0$) and f is linear. Suppose that higher θ indexes more risky distribution of ε in the sense of second-order stochastic dominance. Suppose that the

scalar $F_{13} + F_{11} < 0$ *and that the scalar* $f'(F_{11} + F_{13}) + F_{21}\left(1 + \frac{F_{13}}{F_{11}}\right) < 0$. *Then,* x_0 *is increasing in* θ.

Proof. For a given distribution of ε, i.e. a given θ, the optimal x_0 satisfies the first-order condition

$$E_\varepsilon J_1(x_0^*(\theta), \varepsilon) = 0.$$

It is straightforward from the implicit function theorem that

$$\frac{dx_0^*}{d\theta} = -\frac{\int_{-\infty}^{+\infty} J_1(x_0^*(\theta), \varepsilon)h_\theta(\varepsilon, \theta)d\varepsilon}{\int_{-\infty}^{+\infty} J_{11}(x_0^*(\theta), \varepsilon)h(\varepsilon, \theta)d\varepsilon},$$

and the denominator is negative by the second-order condition. Given that higher θ indexes more risky distribution, the numerator will be positive if the function J_1 is convex in ε; we will prove this which demonstrates our result.

To prove that J_1 is convex in ε, we first calculate J. Define the unconstrained maximum

$$x_1^*(x_0, \varepsilon) = \arg\max_{x_1} F(x_1, x_0, \varepsilon).$$

This maximum is unique since F is quadratic; indeed, x_1 can be written

$$x_1^*(x_0, \varepsilon) = \alpha x_0 + \beta \varepsilon + \gamma,$$

with $\alpha = -\frac{F_{12}}{F_{11}}$ and $\beta = -\frac{F_{13}}{F_{11}}$. We then have the following expression for J :

$$\begin{aligned} J(x_0, \varepsilon) &= F(x_1^*(x_0, \varepsilon), x_0, \varepsilon), \text{ if } x_1^*(x_0, \varepsilon) - f(x_0) \leq \varepsilon, \\ &= F(f(x_0) + \varepsilon, x_0, \varepsilon), \text{ if } x_1^*(x_0, \varepsilon) - f(x_0) > \varepsilon, \end{aligned}$$

and using the envelope theorem we calculate

$$\begin{aligned} J_1(x_0, \varepsilon) &= F_2(x_1^*(x_0, \varepsilon), x_0, \varepsilon), \text{if } x_1^*(x_0, \varepsilon) - f(x_0) \leq \varepsilon, \\ &= f'(x_0)F_1(f(x_0) + \varepsilon, x_0, \varepsilon) + F_2(f(x_0) + \varepsilon, x_0, \varepsilon), \text{ if } x_1^*(x_0, \varepsilon) - f(x_0) > \varepsilon. \end{aligned}$$

Since F is quadratic and f is linear, (and hence x_1 is linear), the two expressions for J_1 are both linear in ε. To determine the convexity of this function simply requires comparing the slopes.[55]

More precisely, given the linearity of x_1 in ε, and our assumption that $\beta = -F_{13}/F_{11} < 1$, there is a threshold value $\bar{\varepsilon}$ such that, if $\varepsilon \geq \bar{\varepsilon}$, we are in the first case (i.e. $x_1^*(x_0, \varepsilon) - f(x_0) \leq \varepsilon$), and if $\varepsilon < \bar{\varepsilon}$, we are in the second case (i.e. $x_1^*(x_0, \varepsilon) - f(x_0) > \varepsilon$).

---

[55]Note that J_1 is continuous in ε since at the boundary between the two expression, $F_1(f(u_0) + \varepsilon, u_0, \varepsilon) = F_1(u_1^*(u_0, \varepsilon), u_0, \varepsilon) = 0$ by optimality of u_1.

The slope of J_1, as a function of ε, is

$$
\begin{aligned}
J_{1\varepsilon} &= F_{21}\frac{\partial x_1}{\partial \varepsilon} + F_{23} = F_{23} - \frac{F_{13}F_{21}}{F_{11}} \text{ for } \varepsilon \geq \bar{\varepsilon}, \\
&= f'F_{11} + f'F_{13} + F_{21} + F_{23} \text{ for } \varepsilon < \bar{\varepsilon}.
\end{aligned}
$$

J_1 is convex provided that its slope is increasing, i.e.

$$
f'\left(F_{11} + F_{13}\right) < -F_{21}\left(1 + \frac{F_{13}}{F_{11}}\right).
$$

■

We now return to our original problem. We first rewrite the choice in terms of inflation. As a reminder, the general Bellman equation is

$$
\begin{aligned}
W_t(\pi_{t-1}, x_{t-1}, \rho, u) &= \min_{\pi_t, x_t, i_t} \frac{1}{2}\left(\pi_t^2 + \lambda x_t^2\right) + \beta E_{\rho'} W_{t+1}(\pi_t, x_t, \rho', u'), \\
s.t. \quad : \quad & \\
x_t &= \delta x_{t-1} - \frac{1}{\sigma}\left(i_t - \rho - \pi_{t-1}\right), \\
\pi_t &= \xi\pi_{t-1} + \kappa x_t + u, \\
i_t &\geq 0.
\end{aligned}
$$

To replace the output gap by inflation in this problem, note that

$$
x_t = \frac{\pi_t - \xi\pi_{t-1} - u_t}{\kappa},
$$

and the ZLB constraint can be rewritten as

$$
x_t \leq \delta x_{t-1} + \frac{\rho + \pi_{t-1}}{\sigma},
$$

or

$$
\pi_t \leq \bar{\pi}_t,
$$

where

$$
\begin{aligned}
\bar{\pi}_t &= \xi\pi_{t-1} + \kappa\delta\left(\frac{\pi_{t-1} - \xi\pi_{t-2} - u_{t-1}}{\kappa}\right) + \frac{\kappa}{\sigma}\left(\rho + \pi_{t-1}\right) + u_t \\
&= \left(\xi + \delta + \frac{\kappa}{\sigma}\right)\pi_{t-1} + \frac{\kappa}{\sigma}\rho - \xi\delta\pi_{t-2} - \delta u_{t-1} + u_t
\end{aligned}
$$

It is thus possible to rewrite the Bellman equation as

$$
W_t(\pi_{t-1}, \bar{\pi}_t, \rho, u) = \min_{\pi_t} \frac{1}{2}\left(\pi_t^2 + \frac{\lambda}{\kappa^2}\left(\pi_t - \xi\pi_{t-1} - u\right)^2\right) + \beta E_{\rho'|\rho} W_{t+1}(\pi_t, \bar{\pi}_{t+1}, \rho', u'),
$$

$$
\begin{aligned}
s.t. \quad : \\
\pi_t &\leq \overline{\pi}_t, \\
\overline{\pi}_{t+1} &= \left(\xi + \delta + \frac{\kappa}{\sigma}\right)\pi_t + \frac{\kappa}{\sigma}\rho' - \xi\delta\pi_{t-1} - \delta u + u'.
\end{aligned}
$$

We can simplify this further given our specific scenario. Given that there is no uncertainty for $t \geq 2$ and that the ZLB constraint does not bind, the value function is simply

$$
W(\pi_{t-1}) = \min_{\pi_t} \frac{1}{2}\left(\pi_t^2 + \frac{\lambda}{\kappa^2}(\pi_t - \xi\pi_{t-1})^2\right) + \beta W(\pi_t).
$$

This value function will of course be quadratic:

$$
W(\pi) = \frac{W}{2}\pi^2.
$$

The value function at time $t = 1$ must take into account that the ZLB may bind. We call this value V:

$$
\begin{aligned}
V(\pi_0, \overline{\pi}_1, u_1) &= \min_{\pi_1} \frac{1}{2}\left(\pi_1^2 + \frac{\lambda}{\kappa^2}(\pi_1 - \xi\pi_0 - u_1)^2\right) + \beta\frac{W}{2}\pi_1^2, \\
s.t. \quad : \quad \pi_1 &\leq \overline{\pi}_1.
\end{aligned}
$$

Finally, the time 0 problem is

$$
\begin{aligned}
U(\pi_{-1}, u_0; \theta) &= \min_{\pi_0} \frac{1}{2}\left(\pi_0^2 + \frac{\lambda}{\kappa^2}(\pi_0 - \xi\pi_{-1} - u_0)^2\right) + \beta E_{\rho_1, u_1} V(\pi_0, \overline{\pi}_1, u_1), \\
s.t. \quad : \quad \overline{\pi}_1 &= \left(\xi + \delta + \frac{\kappa}{\sigma}\right)\pi_0 + \frac{\kappa}{\sigma}\rho_1 - \delta\xi\pi_{-1} - \delta u_0 + u_1,
\end{aligned}
$$

where θ indexes the distribution of either ρ_1^n or u_1. Note that once we have solved for π_0, we can find $x_0 = \frac{\pi_0 - \xi\pi_{-1} - u_0}{\kappa}$ and $i_0 = \rho + \pi_{-1} + \sigma(\delta x_{-1} - x_0)$ immediately. Hence a higher (lower) π_0 implies a higher (lower) x_0 and lower (higher) i_0.

To map our problem in the formulation of the lemma, we first consider the case where the uncertainty is over natural rate shocks (so u_1 is known). In this case, we define

$$
F(\pi_1, \pi_0, \varepsilon) = -\frac{1}{2}\left(\pi_0^2 + \frac{\lambda}{\kappa^2}(\pi_0 - \xi\pi_{-1} - u_0)^2\right) - \frac{1}{2}\left(\pi_1^2 + \frac{\lambda}{\kappa^2}(\pi_1 - \xi\pi_0 - u_1)^2\right) - \beta\frac{W}{2}\pi_1^2,
$$

and

$$
f(\pi_0) = \left(\xi + \delta + \frac{\kappa}{\sigma}\right)\pi_0 - \delta\xi\pi_{-1} - \delta u_0 + u_1.
$$

The problem is then

$$
U(\pi_{-1}, u_0; \theta) = \max_{\pi_0} E_\varepsilon J(\pi_0, \varepsilon),
$$

where

$$J(\pi_0, \varepsilon) = \max_{\pi_1} F(\pi_1, \pi_0, \varepsilon)$$
$$s.t. \quad : \quad \pi_1 \le f(\pi_0) + \varepsilon,$$

with $\varepsilon = \frac{\kappa}{\sigma}\rho_1$. Clearly F is quadratic and f is linear. We have $F_{13} = 0$ so $F_{11} + F_{13} < 0$ is satisfied, and

$$f'(F_{11} + F_{13}) + F_{21}\left(1 + \frac{F_{13}}{F_{11}}\right) = f'F_{11} + F_{12} = -\left(\xi + \delta + \frac{\kappa}{\sigma}\right)(\beta W + 1) - \frac{\lambda}{\kappa^2}\left(\delta + \frac{\kappa}{\sigma}\right) < 0,$$

so the theorem applies, i.e. π_0 (and hence x_0, i_0) is increasing in θ.

To now apply our result in the case of cost-push shocks, we define

$$f(\pi_0) = \left(\xi + \delta + \frac{\kappa}{\sigma}\right)\pi_0 - \delta\xi\pi_{-1} - \delta u_0 + \frac{\kappa}{\sigma}\rho_1,$$

and $\varepsilon = u_1$ (and assume ρ_1 is known). We now need to verify the two conditions. First,

$$F_{13} + F_{11} = \frac{\lambda}{\kappa^2} - \left(\beta W + 1 + \frac{\lambda}{\kappa^2}\right) = -(\beta W + 1) < 0.$$

Second,

$$f'(F_{11} + F_{13}) + F_{21}\left(1 + \frac{F_{13}}{F_{11}}\right),$$
$$= -\left(\xi + \delta + \frac{\kappa}{\sigma}\right)(\beta W + 1) + \frac{\lambda}{\kappa^2}\xi\frac{\beta W + 1}{\beta W + 1 + \frac{\lambda}{\kappa^2}}$$
$$< -\left(\xi + \delta + \frac{\kappa}{\sigma}\right)(\beta W + 1) + \frac{\lambda}{\kappa^2}\xi\frac{\beta W + 1}{\frac{\lambda}{\kappa^2}}$$
$$< -\left(\delta + \frac{\kappa}{\sigma}\right)(\beta W + 1) < 0.$$

D   Numerical methods used for simulations

We present here the numerical methods used to solve the forward looking and the backward looking model. In both cases, we make the following assumptions regarding exogenous variables. First, there is a date T such that, for $t \ge T$, the cost-push shock is zero and the natural rate is constant, $u_t = 0$ and $\rho_t^n = \bar{\rho}$. Second, for $t < T$, the cost-push shock u_t follows a Markov chain with transition probability $P_u(u'|u)$. The natural rate ρ_t^n is the sum of a deterministic component and a Markov chain: $\rho_t^n = f_t + \varepsilon_t$, where ε_t has transition $P_\varepsilon(\varepsilon'|\varepsilon)$, and f_t is increasing and satisfies $f_T = \bar{\rho}$. We will write $\rho_t^n(\varepsilon) = f_t + \varepsilon$. The stochastic processes ε_t and u_t are independent. In practice we use simply $f_t = \rho_0^n + \frac{t}{T_0}(\bar{\rho} - \rho_0^n)$ for $0 \le t \le T_0$

and $f_t = \overline{\rho}$ for $T_0 \leq t < T$. We will choose the Markov chains for ε and for u to each approximate an AR(1) process using the Rouwenhorst method. Our Matlab(r) code will be made available on the following website: `https://sites.google.com/site/fgourio/`

D.1 Forward-looking model solution method

The model we study is

$$
\begin{aligned}
\pi_t &= \beta E_t \pi_{t+1} + \kappa x_t + u_t, \\
x_t &= E_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \rho_t^n(\varepsilon) - E_t \pi_{t+1}\right), \\
i_t &\geq 0.
\end{aligned}
$$

Our theoretical analysis assumed for simplicity (and as is common in the literature) a zero inflation steady-state. To provide more useful numerical illustrations, we consider the case of a positive inflation target. We assume that the equations above apply if π_t is inflation deviation from target and i_t is the nominal rate minus the inflation target. The ZLB is then modified as $i_t \geq \overline{Z} \overset{\text{def}}{=} -\pi^*$.[56]

D.1.1 Calculation of optimal policy under discretion

Optimal policy under discretion can be easily calculated in this model. For $t \geq T$, we have $x_t = \pi_t = 0$. For $t < T$, the optimal policy is given by the solution to

$$
\begin{aligned}
L_t(\varepsilon, u) &= \min_{i_t \geq \overline{Z}} \frac{1}{2}\left(\pi_t^2 + \lambda x_t^2\right) + \beta E_t L_{t+1}(\varepsilon', u') \\
s.t. \quad &: \\
\pi_t &= \beta E_t \pi_{t+1} + \kappa x_t + u, \\
x_t &= E_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \rho_t^n(\varepsilon) - E_t \pi_{t+1}\right),
\end{aligned}
$$

where future expectations $E_t \pi_{t+1}$ and $E_t x_{t+1}$ are taken as given. Since the current decision for i_t does not affect the future loss L_{t+1}, the optimal choice is found by simply minimizing $\pi_t^2 + \lambda x_t^2$.

---

[56]One technical issue is that the long-run Phillips curve is not vertical in this model. To make sure that π^* is indeed the long-run inflation when there is no uncertainty, we assume that the true model is

$$
\pi_t = \beta E_t \pi_{t+1} + (1-\beta)\pi^* + \kappa x_t + u_t,
$$

and the IS curve is unchanged. The policymaker objective is to minimize the expected discounted sum of $(\pi_t - \pi_t^*)^2 + \lambda x_t^2$. We can then redefine $\widetilde{\pi}_t = \pi_t - \pi^*$ and $\widetilde{i}_t = i_t - \pi^*$. The model is now exactly the one written above. Our modification of the Phillips curve is minimal since $(1-\beta)\pi^*$ is a very small number. We make the same assumption in the backward-looking model.

Denote
$$a_t(\varepsilon, u) = E_t\left(x_{t+1} | \varepsilon_t = \varepsilon \text{ and } u_t = u\right),$$

and
$$b_t(\varepsilon, u) = E_t\left(\pi_{t+1} | \varepsilon_t = \varepsilon \text{ and } u_t = u\right),$$

and define $X_t(\varepsilon, u) = 1$ if ZLB binds at time t in state (ε, u), and 0 if not.

Suppose first that the ZLB does not bind; taking first-order conditions then yields

$$
\begin{aligned}
x_t^{nb}(\varepsilon, u) &= -\frac{\kappa}{\lambda + \kappa^2}\left(\beta E_t \pi_{t+1} + u\right) = -\frac{\kappa}{\lambda + \kappa^2}\left(\beta b_t(\varepsilon, u) + u\right), \\
\pi_t^{nb}(\varepsilon, u) &= \frac{\lambda}{\lambda + \kappa^2}\left(\beta E_t \pi_{t+1} + u\right) = \frac{\lambda}{\lambda + \kappa^2}\left(\beta b_t(\varepsilon, u) + u\right), \\
i_t^{nb}(\varepsilon, u) &= \rho_t^n(\varepsilon) + b_t(\varepsilon, u) + \sigma\left(a_t(\varepsilon, u) - x_t(\varepsilon, u)\right).
\end{aligned}
$$

If this solution is feasible, then it is clearly the optimum. If this solution is not feasible, then the optimum is simply to set the nominal interest rate to zero. Hence, the ZLB binds if the nominal interest rate required to implement that solution is negative, i.e.

$$X_t(\varepsilon, u) = 1 \text{ if } \rho_t^n(\varepsilon) + b_t(\varepsilon, u) + \sigma\left(a_t(\varepsilon, u) + \frac{\kappa}{\lambda + \kappa^2}\left(\beta b_t(\varepsilon, u) + u\right)\right) \leq \overline{Z}.$$

In that case, the solution is:

$$
\begin{aligned}
x_t^{zlb}(\varepsilon, u) &= -\frac{\left(\overline{Z} - \rho_t^n(\varepsilon)\right)}{\sigma} + E_t x_{t+1} + \frac{E_t \pi_{t+1}}{\sigma} = -\frac{\left(\overline{Z} - \rho_t^n(\varepsilon)\right)}{\sigma} + a_t(\varepsilon, u) + \frac{b_t(\varepsilon, u)}{\sigma}, \\
\pi_t^{zlb}(\varepsilon, u) &= \kappa\left(-\frac{\left(\overline{Z} - \rho_t^n(\varepsilon)\right)}{\sigma} + a_t(\varepsilon, u) + \frac{b_t(\varepsilon, u)}{\sigma}\right) + \beta b_t(\varepsilon, u) + u, \\
i_t^{zlb}(\varepsilon, u) &= 0.
\end{aligned}
$$

To solve for the optimal path, we only need to know $a_t(\varepsilon, u)$ and $b_t(\varepsilon, u)$. We can solve for these recursively. We have $a_{T-1}(\varepsilon, u) = b_{T-1}(\varepsilon, u) = 0$ for all ε, u, since $x_T = \pi_T = 0$. To update the recursion, we write

$$
\begin{aligned}
a_t(\varepsilon, u) &= E_t\left(x_{t+1} | \varepsilon_t = \varepsilon, u_t = u\right) \\
&= \sum_{\varepsilon', u'} P_\varepsilon(\varepsilon'|\varepsilon) P_u(u'|u)\left(X_{t+1}(\varepsilon', u') x_{t+1}^{zlb}(\varepsilon', u') + (1 - X_{t+1}(\varepsilon', u')) x_{t+1}^{nb}(\varepsilon', u')\right),
\end{aligned}
$$

and

$$b_t(\varepsilon, u) = E_t\left(\pi_{t+1} | \varepsilon_t = \varepsilon, u_t = u\right)$$

$$= \sum_{\varepsilon',u'} P_\varepsilon(\varepsilon'|\varepsilon)P_u(u'|u)\left(X_{t+1}(\varepsilon',u')\pi_{t+1}^{zlb}(\varepsilon',u') + (1 - X_{t+1}(\varepsilon',u'))\,\pi_{t+1}^{nb}(\varepsilon',u')\right).$$

We can then calculate recursively $x_t(\varepsilon, u)$ and $\pi_t(\varepsilon, u)$ for all t, ε, u; consequently we can calculate the loss function $L_t(\varepsilon, u)$ recursively. Start from

$$L_T(\varepsilon, u) = \sum_{t=T}^{\infty} \beta^t\left(\pi_t^2 + \lambda x_t^2\right) = 0,$$

and backwards for $t = 0...T - 1$:

$$L_t(\varepsilon, u) = \pi_t(\varepsilon, u)^2 + \lambda x_t(\varepsilon, u)^2 + \beta \sum_{\varepsilon',u'} P_\varepsilon(\varepsilon'|\varepsilon)P_u(u'|u)L_{t+1}(\varepsilon', u').$$

D.1.2 Calculation of equilibrium under alternative rules

Suppose that the central bank follows the policy:

$$i_t = \max\left(\overline{Z}, g_t(\varepsilon_t) + \phi\pi_t + \gamma x_t\right),$$

where $g_t(\varepsilon_t)$ is a function; for $t \geq T$, $g_t(\varepsilon_t)$ is assumed to be constant, equal to $\overline{g} > \overline{Z}$. This formulation nests the three examples we study in the paper:

- Set i_t equal to the current natural real rate of interest, $i_t = \max\left(\overline{Z}, \rho_t^n(\varepsilon_t)\right)$; this corresponds to $g_t(\varepsilon_t) = \rho_t^n(\varepsilon_t)$, and $\phi = \gamma = 0$;

- A Taylor rule with a constant intercept, $i_t = \max\left(\overline{Z}, \widehat{\rho} + \phi\pi_t + \gamma x_t\right).$

The system of equations to solve is

$$\pi_t = \beta E_t\pi_{t+1} + \kappa x_t + u_t,$$
$$x_t = E_t x_{t+1} - \frac{1}{\sigma}\left(\max\left(\overline{Z}, g_t(\varepsilon_t) + \phi\pi_t + \gamma x_t\right) - \rho_t^n(\varepsilon_t) - E_t\pi_{t+1}\right),$$

and the difficulty is that which formula applies for the interest rate depends on the value of inflation and the output gap, which themselves depend on the interest rate. However it is easy to solve the model by backward induction, in a way roughly similar to the optimal policy calculation above. For $t \geq T$, $i_t = \overline{g} > 0$, and the equilibrium is

$$\overline{\pi} = \overline{g} - \overline{\rho},$$
$$\overline{x} = \frac{1 - \beta}{\kappa}\overline{\pi}.$$

In particular, if $\overline{g} = \overline{\rho}$, the terminal state is $x = \pi = 0$. (This case is the outcome for cases (a) and (b) but not necessarily for case (c), depending on whether $\widehat{\rho} = \overline{\rho}$.)

Use a superscript W to denote the outcomes with this rule. Define again

$$
\begin{aligned}
a_t^W &= E_t\left(x_{t+1}^W | \varepsilon_t = \varepsilon \text{ and } u_t = u\right), \\
b_t^W &= E_t\left(\pi_{t+1}^W | \varepsilon_t = \varepsilon \text{ and } u_t = u\right),
\end{aligned}
$$

and note that

$$
\begin{aligned}
\pi_t^W(\varepsilon, u) &= \beta b_t^W(\varepsilon, u) + u + \kappa x_t^W(\varepsilon, u) \\
x_t^W(\varepsilon, u) &= a_t^W(\varepsilon, u) - \frac{1}{\sigma}\left(\max\left(\overline{Z}, g_t(\varepsilon_t) + \phi \pi_t^W(\varepsilon, u) + \gamma x_t^W(\varepsilon, u)\right) - \rho_t^n(\varepsilon_t) - b_t^W(\varepsilon, u)\right).
\end{aligned}
$$

There are two possible cases, depending on whether $g_t(\varepsilon_t) + \phi \pi_t^W(\varepsilon, u) + \gamma x_t^W(\varepsilon, u) > \overline{Z}$. Consider first the case where it is positive. In this case, simple algebra yields

$$
x_t^W(\varepsilon, u) = \frac{1}{1 + \frac{\gamma + \kappa\phi}{\sigma}}\left(a_t^W(\varepsilon, u) - \frac{\phi}{\sigma}\left(\beta b_t^W(\varepsilon, u) + u\right) - \frac{1}{\sigma}\left(g_t(\varepsilon_t) - \rho_t^n(\varepsilon_t) - b_t^W(\varepsilon, u)\right)\right),
$$

and $\pi_t^W(\varepsilon, u)$ can be obtained from the equation above. We can now check if indeed $g_t(\varepsilon_t) + \phi \pi_t^W(\varepsilon, u) + \gamma x_t^W(\varepsilon, u) > \overline{Z}$ is satisfied. If it is not, we then look for a solution at the ZLB, i.e.

$$
\begin{aligned}
\pi_t^W(\varepsilon, u) &= \beta b_t^W(\varepsilon, u) + u + \kappa x_t^W(\varepsilon, u) \\
x_t^W(\varepsilon, u) &= a_t^W(\varepsilon, u) - \frac{1}{\sigma}\left(\overline{Z} - \rho_t^n(\varepsilon_t)\right) + \frac{1}{\sigma}b_t^W(\varepsilon, u),
\end{aligned}
$$

and we check that with this solution, $g_t(\varepsilon_t) + \phi \pi_t^W(\varepsilon, u) + \gamma x_t^W(\varepsilon, u) < \overline{Z}$.[57] Given the value of $\pi_t^W(\varepsilon, u)$ and $x_t^W(\varepsilon, u)$ for all ε, u, we can update $a_{t-1}^W(\varepsilon, u)$ and $b_{t-1}^W(\varepsilon, u)$ and hence proceed backwards until time 0. We can furthermore calculate the loss function in the same way as for optimal policy.

D.2 Backward-looking model solution method

The model is

$$
\pi_t = \xi \pi_{t-1} + \kappa x_t + u_t,
$$

and

$$
x_t = \delta x_{t-1} - \frac{1}{\sigma}\left(i_t - \rho_t^n(\varepsilon_t) - \pi_{t-1}\right).
$$

As in the forward-looking case, we assume that this model applies to deviations of inflation from the target, and i_t is the difference between the nominal rate and the inflation target π^*. (Formally, this can be justified as in the previous footnote.) As a result, we have the ZLB constraint $i_t \geq \overline{Z} = -\pi^*$.

---

[57]In principle, it is possible that either none, or both solutions exist, but we never encountered this case in our calculations.

D.2.1 Calculation of optimal policy under discretion

The optimal policy under discretion can be set up using a Bellman equation:

$$
V_t(x_{-1}, \pi_{-1}, \varepsilon, u) = \min_{i \geq \overline{Z}} \frac{1}{2} \left(\pi^2 + \lambda x^2 \right) + \beta E_{\varepsilon', u' | \varepsilon, u} V_{t+1}(x, \pi, \varepsilon', u'),
$$

$$
s.t. \quad :
$$

$$
\pi = \xi \pi_{-1} + \kappa x + u,
$$

$$
x = \delta x_{-1} - \frac{1}{\sigma} \left(i - \rho_t^n(\varepsilon) - \pi_{-1} \right).
$$

We first solve for the value in the final steady-state, $V_T(x_{-1}, \pi_{-1})$; we have a closed form solution if the ZLB does not bind for all values of x, π (see appendix B); or we can solve it numerically using the Bellman equation

$$
V_T(x_{-1}, \pi_{-1}) = \min_{i \geq 0} \frac{1}{2} \left(\pi^2 + \lambda x^2 \right) + \beta V_T(x, \pi),
$$

$$
s.t. \quad :
$$

$$
\pi = \xi \pi_{-1} + \kappa x,
$$

$$
x = \delta x_{-1} - \frac{1}{\sigma} \left(i - \overline{\rho} - \pi_{-1} \right).
$$

For $t < T$, we solve numerically the Bellman equation above. For simplicity, we assume that only a discrete set of interest rates is allowed, call it $G = \{i_1, ..., i_N\}$. We then solve this Bellman equation by interpolating the value functions around a grid for x and for π. Specifically, at time t, and for each value of x and π in these grids, we calculate the payoff of using any given interest rate $i \in G$ today, and select the optimal one. This may require us to interpolate to find the expected future value; we use a linear interpolation. This solution method produces the optimal policy $i_t(x_{-1}, \pi_{-1}, \varepsilon, u)$ and the output gap and inflation $x_t(x_{-1}, \pi_{-1}, \varepsilon, u)$ and inflation $\pi_t(x_{-1}, \pi_{-1}, \varepsilon, u)$ as well as the loss function $V_t(x_{-1}, \pi_{-1}, \varepsilon, u)$ for all points in the grid. We then move to on to period $t - 1$, and so on until time 0.

D.2.2 Equilibrium under a different policy rule

We can also calculate the equilibrium in this model under a rule of the form

$$
i_t = \max \left(0, g_t(\varepsilon_t) + \phi \pi_t + \gamma x_t \right).
$$

Specifically, given x_{t-1} and π_{t-1} and the values of ε_t, u_t, we must solve the system:

$$
\pi_t = \xi \pi_{t-1} + \kappa x_t + u_t,
$$

$$
x_t = \delta x_{t-1} - \frac{1}{\sigma} \left(\max \left(0, g_t(\varepsilon_t, u_t) + \phi \pi_t + \gamma x_t \right) - \rho_t^n(\varepsilon_t) - \pi_{t-1} \right),
$$

and so we need to consider the two possible cases to find the solution. Either $g_t(\varepsilon_t) + \phi\pi_t + \gamma x_t > 0$, in which case

$$x_t = \frac{1}{1 + \frac{\gamma}{\sigma} + \kappa\frac{\phi}{\sigma}} \left(-\frac{\phi}{\sigma}\xi\pi_{t-1} - \frac{\phi}{\sigma}u_t + \delta x_{t-1} - \frac{1}{\sigma}\left(g_t(\varepsilon_t, u_t) - \rho_t^n(\varepsilon_t) - \pi_{t-1}\right)\right),$$

and $\pi_t = \xi\pi_{t-1} + \kappa x_t + u_t$; and we need to verify that indeed $g_t(\varepsilon_t) + \phi\pi_t + \gamma x_t > 0$; or we have

$$x_t = \delta x_{t-1} - \frac{1}{\sigma}\left(-\rho_t^n(\varepsilon_t) - \pi_{t-1}\right),$$

$$\pi_t = \xi\pi_{t-1} + \kappa x_t + u_t,$$

and we need to verify that indeed $g_t(\varepsilon_t) + \phi\pi_t + \gamma x_t < 0$.

D.2.3 Deflationary traps

An important issue in this model is the risk of deflation trap (Reifschneider and Williams (2000)). For given parameters, there is a set of initial values x_{-1} and π_{-1} that diverges to $-\infty$ even under optimal policy, at least for some shock realizations. Mechanically, this arises because if the output gap is negative and ξ is large enough, inflation will fall; and the output gap will likely fall is δ is large enough and/or the natural rate or inflation are negative enough. Hence, low inflation and output gap can be self-reinforcing. These deflation traps capture an economically meaningful mechanism, but obviously the divergence to $-\infty$ is extreme. In reality, it seems more likely that the divergence would stop at some point due to a regime change in the way policy, expectations, or price setting is determined. For instance, fiscal policy might step in at some point and ensure that the deflation does not perpertuate itself. In our solution method, we impose this - i.e. there is a worst possible outcome, $\underline{\pi}$ for inflation and $\underline{x}$ for the output gap, which "caps" inflation and output gap and hence prevents the divergence to $-\infty$. Obviously, our simulations start from initial conditions such that the deflation trap can be avoided by appropriate policy, so the assumptions regarding the deflation trap are not key to our results. However, policy in this model is also motivated by the desire to prevent the economy from falling under a deflation trap should a negative sequence of shocks arise, and for some parameters this can have a significant effect to increase the "buffer stock" i.e. stay with inflation and output gap above target persistently.