CHARLES EVANS
Federal Reserve Bank of Chicago

FRANÇOIS GOURIO
Federal Reserve Bank of Chicago

JONAS FISHER
Federal Reserve Bank of Chicago

SPENCER KRANE
Federal Reserve Bank of Chicago

Risk Management for Monetary Policy Near the Zero Lower Bound

ABSTRACT With projections showing 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 many standard models uncertainty has no effect. In this paper, we demonstrate that the zero lower bound (ZLB) 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 demonstrate this result theoretically through two canonical macroeconomic models. Using numerical simulations of our models calibrated to the current environment, we find that optimal policy calls for a delay in liftoff of two to three quarters relative to a policy that does not take into account uncertainty about policy being constrained by the ZLB. We then use a narrative study of Federal Reserve communications and estimated policy reaction functions to show that risk management is a long-standing practice in the conduct of monetary policy.

To what extent should uncertainty affect monetary policy? This classic question is relevant today as the Federal Reserve considers when to start increasing the federal funds rate. In the March 2015 "Summary of Economic Projections," 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, returning to its 2 percent target. This forecast could go wrong in two ways. First, the FOMC may be overestimating the underlying strength in the economy or the tendency of inflation to return to target. Guarding against these risks would call for cautious removal of accommodation. Second, the economy could be poised for stronger growth and inflation than currently projected. This second risk would call for more aggressive rate hikes. How should policy manage these divergent risks?

If the FOMC misjudges the impediments to growth and inflation and reduces monetary accommodation too soon, it could find itself in the uncomfortable position of having to reverse course and being constrained by the zero lower bound (ZLB) again. It is true that the FOMC has access to unconventional policy tools at the ZLB, but these appear to be imperfect substitutes for the traditional funds rate instrument. In contrast, if the Fed keeps rates too low and inflation rises too quickly, most likely inflation could be brought back into check with modest increases in interest rates. Since the unconventional tools available to counter the first scenario may be less effective than the traditional tools available to counter the second scenario, the costs of premature liftoff may exceed those of delay. It therefore seems prudent to refrain from raising rates until the FOMC is highly certain that growth is sustainable and inflation is returning to target.[1]

In this paper we establish theoretically that uncertainty about monetary policy being constrained by the ZLB in the future implies an optimally looser policy today, which in the current context means delaying liftoff—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, using standard models of monetary policy, that 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 dictates some counteracting policy easing today. The second channel, which we call the *buffer stock channel,* arises because, if inflation or output is intrinsically persistent, building up output or inflation today reduces the likelihood and severity of hitting the ZLB tomorrow. Optimal policy when either of these channels is operative should be looser whenever a return to the ZLB remains a distinct possibility. In simulations calibrated to the current environment, we find that optimal policy prescribes two to three quarters of delay in liftoff relative to a policy that does not take this uncertainty into account. However, under the

1. In his speech at the Petersen Institute for Economics, Evans (2014) discussed these issues at greater length.

optimal policy the central bank must also be prepared to raise rates quickly as the threat of being constrained by the ZLB recedes.

Would it be unusual for the Fed to take uncertainty into account in setting its policy rate? The second part of this paper argues that risk management has been a long-standing practice in U.S. monetary policy. Therefore, advocating it in the current policy environment would be consistent with a well-established approach of the Federal Reserve. Of course, because the ZLB was only recently perceived as an important constraint, the theoretical rationales for risk management were different in the past. It is true that in a wide class of models that abstract from the ZLB, optimal policy involves adjusting the interest rate in response to the mean of the distribution of shocks, and information on higher moments is irrelevant (the so-called "certainty equivalence" principle). However, there 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.

We explore whether policymakers actually practiced risk management prior to the ZLB period in two ways. First, we analyze Federal Reserve communications over the period 1987–2008 and find numerous examples when uncertainty or the desire to insure against important risks to the economy were used to help explain the setting of policy. Confirmation of this view is found in the statements of Alan Greenspan, who during his tenure as Federal Reserve chair noted, "the conduct of monetary policy in the United States has come to involve, at its core, crucial elements of risk management."[2] Second, we estimate a conventional forecast-based monetary policy reaction function augmented with a variety of measures of risk based on financial market data, Federal Reserve Board staff forecasts, private-sector forecasts, and narrative analysis of the FOMC minutes. We find clear evidence that when measured in this way, risk has had a statistically and economically significant impact on the interest rate choices of the FOMC. For the FOMC, risk management appears to be old hat.

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. We do not think this is the case. Even though most central bankers believe unconventional policies such as large-scale asset purchases (LSAPs) or

2. From "Risk and Uncertainty in Monetary Policy," Chairman Greenspan's remarks at the Meeting of the American Economic Association, San Diego, California, January 3, 2004 (Greenspan 2004).

more explicit and longer-term forward guidance about policy rates can provide considerable accommodation at the ZLB, few argue that these tools are on an equal footing with traditional policy instruments.[3]

One reason for this is that the effects on the economy of unconventional policies are, naturally, much more uncertain than those of traditional tools. There are divergent empirical estimates of their effects, and there is uncertainty about the theoretical mechanism behind those effects. Various studies of LSAPs, for example, provide a wide range of estimates of their ability to put downward pressure on private borrowing rates and influence the real economy. Furthermore, the effects on interest rates of both LSAPs and forward guidance are complicated functions of private-sector expectations, which make their economic effects highly uncertain as well.[4]

Uncertainty about the transmission mechanism of LSAPs is reflected in Arvind Krishnamurthy and Annette Vissing-Jørgensen's (2013) discussion of the various hypotheses that have been proposed. Unconventional tools also carry potential costs. The four most commonly cited costs are these: (i) the large increases in reserves generated by LSAPs risk unleashing inflation; (ii) a large balance sheet may make it more difficult for the Fed to raise interest rates when the time comes; (iii) the extended period of very 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 (iv) 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.[5] Costs reduce

3. For example, while there is econometric evidence that changes in term premia influence activity and inflation, some studies find that 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).

4. Bomfin and Meyer (2010), D'Amico and King (2013), and Gagnon, Raskin, Remache, and Sack (2010) find noticeable effects of LSAPs on Treasury term premia, while Chen, Curida, and Ferrero (2012) and Hamilton and Wu (2010) unearth only small effects. Krishnamurthy and Vissing-Jørgensen (2013) argue that LSAPs only had a substantial influence on private borrowing rates in the mortgage market. Engen, Laubach, and Reifschneider (2015) and Campbell and others (2012) analyze the interactions between LSAPs' forward guidance and private sector expectations.

5. These costs are mitigated, however, by additional tools the Fed has introduced to exert 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 instability concerns. Furthermore, continued low rates of inflation and contained private-sector inflationary expectations have reduced concerns regarding an outbreak of inflation.

the incentive to use any policy tool. Moreover, because the costs of unconventional tools are very hard to quantify, the level of uncertainty associated with them is naturally elevated as well.

A consequence of this uncertainty over the benefits and costs of unconventional tools is that they are likely to be used more cautiously than traditional policy instruments, as suggested by William Brainard's (1967) classic analysis. For example, then Federal Reserve Chairman Ben Bernanke emphasized in 2012 that because of their uncertain costs and benefits, "the hurdle for using nontraditional policies should be higher than for traditional policies."[6] In addition, some of the benefits of unconventional policies may be decreasing, and their costs may be increasing in terms of balance sheet size or amount of time spent in a very low interest rate environment.[7] Accordingly, policies that had widespread support early on in a ZLB episode might be difficult to extend or expand with an already large balance sheet.

So, while they can be valuable, unconventional 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 they face during more conventional times, and thus they are worthy of consideration in their own right. We abstract from unconventional policy tools for the remainder of our analysis.

I. 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. 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

6. "Monetary Policy since the Onset of the Crisis," Remarks by Chairman Ben S. Bernanke at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, August 31, 2012 (Bernanke 2012, p. 14).

7. Krishnamurthy and Vissing-Jørgensen (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 were forced to expand their use in light of unexpected economic weakness.

or inflation.[8] 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 sections II and III suggests that in practice, policymakers are sensitive to uncertainty and respond to it by following what appears to be a risk-management approach. Understanding 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. First, we use the workhorse forward-looking New Keynesian model to illustrate the expectations channel, in which the possibility of a binding ZLB tomorrow leads to lower expected inflation and an output gap occurring today, thus necessitating policy easing today. We then use a backward-looking "Old" Keynesian setup 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. Both of these channels operate in modern DSGE (dynamic stochastic general equilibrium) models such as those described by Lawrence Christiano, Martin Eichenbaum, and Charles Evans (2005) and by Frank Smets and Rafael Wouters (2007), but they are more transparent if we consider them in separate, although related, simple models. After describing these two channels we construct some numerical simulations to assess their quantitative effects.

I.A. The Expectations Channel

The simple New Keynesian model has well established micro-foundations based on price stickiness. Given that excellent expositions of these foundations have been offered many times, for example by Michael Woodford (2003) and Jordi Galí (2008), we simply 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

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

8. 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. The Federal Reserve Board staff routinely conducts optimal policy exercises in the FRB/US model; see for example English, Lopez-Salido, and Tetlow (2013).

where 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 with rational expectations is assumed; u_t is a mean zero exogenous *cost-push* shock; and $0 < \beta < 1$, $\kappa > 0$. 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. 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

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

The variable g_t is an exogenous mean zero *demand* shock, and z_t is the exogenous log of potential output. Since g_t and z_t are exogenous, so is the natural rate. Equation 2 indicates that ρ_t^n corresponds to the setting of the nominal interest rate consistent with expected inflation at target and the output gap equal to zero.[9] If potential output is constant and the demand shock equals zero, then the natural rate equals the constant $\bar{\rho} > 0$.

Our analysis is centered on uncertainty in the natural rate.[10] From equation 3 we see that this uncertainty derives from 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 deleveraging.[11] The latter source of uncertainty is over the variety of factors that can influence the expected rate of growth in potential

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. There is ample evidence of considerable uncertainty regarding the natural rate. See for example Barsky, Justiniano, and Melosi (2014), Hamilton and others (2015), and Laubach and Williams (2003).

11. 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 in turn depresses demand, as is emphasized by Basu and Bundick (2014), Fernandez-Villaverde and others (2012), and Born and Pfeifer (2014).

output, for example as emphasized in the recent debate over so-called secular stagnation.

We adopt the canonical framework in assuming that the central bank acts to minimize a quadratic loss function with the understanding that private-sector behavior is governed by equations 1 through 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, that is, $i_t \geq 0$, abstracting from the possibility that the effective lower bound on i_t is slightly negative. The short-term interest rate is the central bank's only policy instrument, and it is set by solving for optimal policy under discretion. In particular, in each period the central bank sets the nominal interest rate with the understanding that private agents anticipate that it will re-optimize in the following periods.

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 Paul Krugman (1998), Gauti Egertsson and Michael Woodford (2003), Woodford (2012), and Iván 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.[12] Nevertheless, below we discuss our intuition for why our main result should extend to the case of commitment. Second, discretion may better approximate the institutional environment in which the FOMC operates.

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;

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

moreover, there is no uncertainty in the natural rate after $t = 2$, $\rho_t^n = \overline{\rho} > 0$ 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.[13] The variable ρ_1^n is assumed to be distributed according to the probability density function $f_\rho(\cdot)$.

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; notably, our formulation allows us to consider the optimal timing of liftoff. We do not have to take a stand on whether the ZLB is binding before $t = 0$, but one possibility is that the natural rate ρ_t^n was sufficiently negative for $t < 0$ so that the optimal policy rate was set at zero, $i_t = 0$ for $t < 0$, but because the economy has been improving the natural rate is close to zero by $t = 0$. The question is whether to raise the policy rate at $t = 0$, $t = 1$, or $t = 2$.

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, $i_t = \rho_t^n = \overline{\rho}$. 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 $x_1 = \rho_1^n/\sigma < 0$. The expected output gap at $t = 1$ is $E_0 x_1 = \int_{-\infty}^{0} \rho f_\rho(\rho) d\rho \big/ \sigma \leq 0$ 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. Low inflation tomorrow depresses inflation today, since price-setting is forward-looking in the Phillips curve, and it also depresses output today by raising the real interest rate via the IS curve. The optimal policy at $t = 0$ must take into account these effects. This implies that optimal policy will be looser than if there were no chance that the ZLB would bind tomorrow.

13. 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 the amount of uncertainty would have no effect. This is why our scenario has more than two periods.

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 Galí (2008, pp. 76–77) for a discussion. Henceforth we will not consider this issue.

Mathematically, substituting for π_0 and i_0 using equations 1 and 2, and taking into account the ZLB constraint, optimal policy at $t = 0$ solves the following problem:

$$(5) \qquad \min_{x_0} \frac{1}{2}\left[\left(\kappa x_0 + \beta E_0 \pi_1 + u_0\right)^2 + \lambda x_0^2\right] \text{ s.t. } x_0 \leq E_0 x_1 + \frac{1}{\sigma}\left(\rho_0^n + E_0 \pi_1\right).$$

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

$$(6) \qquad \rho_0^* = -\sigma \frac{\kappa}{\lambda + \kappa^2} u_0 - \left(1 + \frac{\kappa}{\sigma} + \beta \frac{\kappa^2}{\lambda + \kappa^2}\right) \int_{-\infty}^{0} \rho f_\rho(\rho) d\rho.$$

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

$$(7) \qquad x_0 = -\frac{\kappa}{\lambda + \kappa^2}\left(\beta E_0 \pi_1 + u_0\right); \quad \pi_0 = \frac{\lambda}{\lambda + \kappa^2}\left(\beta E_0 \pi_1 + u_0\right).$$

The corresponding interest rate is

$$(8) \qquad \begin{aligned} i_0 &= \rho_0^n + E_0 \pi_1 + \sigma\left(E_0 x_1 - x_0\right) \\ &= \rho_0^n + \sigma \frac{\kappa}{\lambda + \kappa^2} u_0 + \left(1 + \frac{\kappa}{\sigma} + \beta \frac{\kappa^2}{\lambda + \kappa^2}\right) \int_{-\infty}^{0} \rho f_\rho(\rho) d\rho. \end{aligned}$$

As long as $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho < 0$, equation 8 implies that the optimal interest rate is lower than if there were no chance of a binding ZLB tomorrow, that is, if $f_\rho(\rho) = 0$ 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 $i_0 = 0$. In this case,

$$(9) \qquad x_0 = \frac{\rho_0^n}{\sigma} + \left(1 + \frac{\kappa}{\sigma}\right) E_0 x_1; \quad \pi_0 = \kappa \frac{\rho_0^n}{\sigma} + \left[(1 + \beta)\kappa + \frac{\kappa^2}{\sigma}\right] E_0 x_1.$$

Notice from equation 6 that higher uncertainty makes it more likely that the ZLB will bind at $t = 0$. Specifically, even 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 of being constrained by the ZLB today.

Since $E_0 x_1$ is a sufficient statistic for $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho$ in equation 8, the optimal policy has the flavor of a traditional forward-looking policy reaction function that only depends on the conditional expectations of output and inflation gaps. However $E_0 x_1$ is not independent of a mean-preserving spread or any other change in the distribution of ρ_1^n. Accordingly, optimal policy here departs from the certainty equivalence principle, which says that the extent of uncertainty in the *underlying fundamentals* (in our case ρ_1^n) does not affect the optimal interest rate.[15] Furthermore, as a practical matter the central bank must infer private agents' $E_0 x_1$ in order to determine optimal policy. Since $E_0 x_1$ depends on the entire distribution of ρ_1^n, so must the central bank's estimates of it, which is a much more difficult inference problem than in the certainty equivalence case.

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

Proposition 1: *Higher uncertainty, that is, a mean-preserving spread in the distribution of the natural rate* ρ_1^n *tomorrow, leads to a looser optimal 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, that is, more negative, $E_0 x_1$ and $E_0 \pi_1$. Hence, higher uncertainty leads to lower i_0.[16]

The effect of higher uncertainty on i_0 is unambiguous, but the effect on the output gap and inflation is more subtle. If the ZLB does not bind at $t = 0$ initially, higher uncertainty leads to lower $E_0 x_1$ and $E_0 \pi_1$ and consequently to higher x_0 and lower π_0 according to equation 7. On the other hand, if the ZLB *does* bind at $t = 0$ initially, then higher uncertainty leads to lower x_0 and lower π_0 according to equation 9.[17] Overall, the effect of higher uncertainty on π_0 is unambiguously negative, but the effect on x_0 may be positive or negative.

15. Recent statements of the certainty equivalence principle in models with forward-looking variables can be found in Svensson and Woodford (2002, 2003).

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

17. Finally, there is a case where the ZLB does not bind initially but does bind if uncertainty is higher. In this case, x_0 may be lower or higher with higher uncertainty, while π_0 is always smaller.

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 today is adjusted only with respect to the states of the 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, natural rate shocks can and should be perfectly offset by monetary policy.

DISCUSSION Proposition 1 has several predecessors. Perhaps the closest are Klaus Adam and Roberto Billi (2007), Taisuke Nakata (2013a,b), and Anton 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 lower 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 it were sure the ZLB constraint would never bind again.[19] Even though the best forecast may be that the economy will recover and exit the ZLB—that is, 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; that is, when $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho$ is a large negative number; the economy is less sensitive to interest rates (high σ); and the Phillips curve is steep (high κ).

With higher uncertainty, the increase in interest rates will be faster on average from $t = 0$ to $t = 2$. This follows since the $t = 2$ interest rate is unaffected by uncertainty whereas at $t = 0$ it is lower. More generally, when uncertainty about being constrained by the ZLB in the future dissipates, the interest rate can rise quickly because the effects holding it down disappear along with the uncertainty.

While we have deliberately focused on a very simple example, our results hold under 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 online appendix we consider the case of optimal policy

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

19. Indeed, private sector forecasters attribute a significant likelihood of a return to the ZLB: respondents to the January 2015 Federal Reserve Bank of New York survey of Primary Dealers put the odds of returning to the ZLB within two years following liftoff at 20 percent.

with uncertainty about cost-push inflation.[20] We show that optimal policy also is looser if there is a chance of a binding ZLB in the future due to a low cost-push shock. Furthermore, 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 would 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.

Intuitively, we expect a version of Proposition 1 to still hold with commitment as well. Optimal policy with commitment involves promising at $t = 0$ that should the ZLB bind at $t = 1$, the central bank would keep interest rates lower for $t \geq 2$ than it would otherwise. As is well known, this policy reduces the size of the inflation and output gaps at $t = 1$, but it does not eliminate them entirely. These gaps then could generate negative expected inflation and output gaps at $t = 0$ that become more negative the larger the $t = 1$ uncertainty. Higher uncertainty should therefore lead to looser policy at $t = 0$, just as in the case of discretion.

One obvious limitation to these results is that we have assumed (and will continue do so when studying the backward-looking model below) that there is no cost to raising rates quickly if needed. For example, our welfare criterion does not value interest-rate smoothing. Smoothing has been rationalized by Marvin Goodfriend (1991) and others as facilitating financial market adjustments or as a signaling tool. It is true also that estimated reaction functions include lagged funds rate terms to fit historical data. Nonetheless, there have been instances when the FOMC has moved quickly. Some of these occurred as recessions unfolded, but not all: between February 1994 and February 1995 rates were tightened by 300 basis points and between November 1988 and February 1989 by nearly 165 basis points. Moreover, as Brian Sack (2000) and Glenn Rudebusch (2002) argue, interest rate smoothing might reflect learning about an uncertain economy rather than a desire to avoid large changes in interest rates per se. The policy prescriptions derived from our models are specifically aimed at addressing such uncertainty.

I.B. The Buffer Stock Channel

The buffer stock channel relies not on forward-looking behavior but on the view that the economy has some inherent *momentum,* for instance due

20. Online appendixes to all papers in this volume may be found at the BPEA web page, www.brookings.edu/bpea, under "Past Editions."

to adaptive inflation expectations, inflation indexation, habit persistence, adjustment costs, or hysteresis. Suppose that output or inflation has 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 be weighed against the costs of excess output and inflation today, as well as tomorrow's cost to bring down the output gap or inflation if the ZLB turns out not to bind. 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 as the New Keynesian model does, but backward-looking elements appear to be important empirically.[21] Backward-looking models have been studied extensively in the literature, including by Thomas Laubach and John Williams (2003), Athanasios Orphanides and Williams (2002), David Reifschneider and Williams (2000), and Rudebusch and Lars Svensson (1999).

The model we study simply replaces the forward-looking terms in equations 1 and 2 with backward-looking terms:

$$\pi_t = \xi\pi_{t-1} + \kappa x_t + u_t; \tag{10}$$

$$x_t = \delta x_{t-1} - \frac{1}{\sigma}\left(i_t - \rho_t^n - \pi_{t-1}\right), \tag{11}$$

where $0 < \xi < 1$ and $0 < \delta < 1$. This model is essentially the same as the simple example Reifschneider and Williams (2000) use to motivate their analysis of monetary policy constrained by the ZLB. 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 equation 3 as a guide to

21. Indeed, empirical studies based on medium-scale DSGE models, such as those considered by Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007), find that 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), Galí and Gertler (1999), and Eichenbaum and Fisher (2007).

interpreting it, but it is perhaps better to think of it as simply a "demand" shock or "IS" shock.

ANALYSIS We consider the ZLB scenario described in section I.A (under "A ZLB Scenario") and again 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, 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 $t = 2$, that is, if $\rho_2^n = \overline{\rho}$ is large enough, an assumption we will maintain.[22] Given the quadratic loss, it is optimal to smooth this adjustment over time so that the economy will converge back to its steady-state slowly. The details of this adjustment after $t = 2$ are not very important for our analysis. What is important is that the overall loss of starting from $t = 2$ with lagged inflation π_1 and output gap x_1 is 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 online 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 equations 10 and 11, and taking into account the ZLB constraint, optimal policy at $t = 1$ solves the following problem:

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

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:

$$\rho_1^*(x_0, \pi_0) = -\left[\frac{(1+\beta W)\kappa\xi}{(1+\beta W)\kappa^2 + \lambda}\sigma + 1\right]\pi_0 - \sigma\delta x_0. \tag{12}$$

For $\rho_1^n \ge \rho_1^*(x_0, \pi_0)$ the ZLB is not binding; otherwise it is. Hence the probability of hitting the ZLB is $\int_{-\infty}^{\rho_1^*(x_0,\pi_0)} f_\rho(\rho)d\rho$. In contrast to the forward-looking case, the probability of being constrained by the ZLB constraint is now endogenous at $t = 1$ and can be influenced by policy at $t = 0$. As

22. Relaxing it would only strengthen our results.

indicated by equation 12, 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

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

This is similar to the forward-looking model's solution, which 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 $\rho_I^n < \rho^*(x_0, \pi_0)$ the ZLB binds, in which case

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

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

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

This expression reveals that the initial conditions x_0 and π_0 matter by shifting the payoff from continuation in the non-ZLB states, $W\pi_0^2/2$; the payoff in the case where the ZLB binds (the second integral); and 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:

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

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* ρ_1^n *tomorrow leads to a looser optimal policy today.*

From equations 10 and 11, higher uncertainty also leads to larger x_0 and π_0. 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.

Our model also implies that an increase in uncertainty over the initial output gap will lead to looser policy. Specifically we have:

Proposition 3: *Suppose the initial output gap* x_{-1} *is unknown at* $t = 0$ *but becomes known at* $t = 1$ *and the central bank has a prior distribution over* x_{-1}*. Then a mean-preserving spread in this prior distribution leads optimal policy to be looser at* $t = 0$*.*

The proof of this proposition is similar to the one for Proposition 2. This result is particularly germane to the current policy environment where there is uncertainty over the amount of slack in the economy. Therefore, Proposition 3 provides an additional rationale for delaying liftoff.

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, liftoff 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, even if not constrained by the ZLB today, an increase in uncertainty about the likelihood of being constrained by the ZLB tomorrow leads to a reduction in the policy rate today. So the buffer stock channel and the expectations channel have very similar policy implications, though 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. Note that as in the forward-looking model, optimal policy prescribes that interest rates rise as the likelihood of being constrained by the ZLB in the future falls, even if the output gap or inflation does not change.

It is useful to compare the policy implications of the buffer stock channel to the argument developed in Olivier Coibion, Yuriy Gorodnichenko, and

Johannes Wieland (2012). Their paper studies the tradeoff between the level of the inflation target and the risk of hitting the ZLB using policy reaction functions instead of optimal policy.[23] Our analysis does not require a drastic change in monetary policy in order to improve outcomes. It is achieved through standard interest-rate policy rather than a credibility-damaging change to the inflation target.

I.C. Quantitative Assessment

We now assess the quantitative significance of the expectations and buffer stock channels using calibrated versions of the forward- and backward-looking models that we solve numerically. With parameters drawn from the literature and initial conditions calibrated to early 2015, we compare equilibrium outcomes under optimal discretion to alternative policies that do not take into account uncertainty. Our numerical methods are described in the online appendix. Importantly, and in contrast to most of the literature, they allow uncertainty to affect policy and to be reflected in welfare.

PARAMETER VALUES The parameter values are reported in table 1. We use the same values for parameters that are common to both models. The time period is one quarter, with $t = 1$ taken to be 2015Q1. The natural rate ρ_t^n is the sum of deterministic and random components. We assume the deterministic component rises linearly between $t = 1$ and $t = T > 1$, after which it remains constant at $\bar{\rho} = 1.75$ percent, which corresponds with the median long-run funds rate in the March 2015 FOMC Summary of Economic Projections, less the FOMC's inflation target $\pi^* = 2$. The random component is AR(1) with auto-correlation coefficient ρ_ε and innovation standard deviation σ_ε. We also assume there is an i.i.d. cost-push shock with standard deviation σ_u. There is no uncertainty for $t > T$.

The degree of uncertainty we assume is central to our findings. The particular values of ρ_ε and σ_ε are not as important to our results as the unconditional volatility they imply. There is wide variation in estimates of volatility in the natural rate, corresponding to differences in theoretical concepts, models and empirical methods used. Our calibration implies that the unconditional standard deviation of the natural rate is 2.5 percent at an annual rate. This lies within the range of estimates in Robert Barsky, Alejandro Justiniano, and Leonardo Melosi (2014), Vasco Cúrdia and others (2015), and Laubach and Williams (2003). The auto-correlation

23. Another difference is that they study a medium-scale DSGE model with both forward- and backward-looking elements; because of this added complexity, they use a different solution method.

Table 1. Parameter Values[a]

Parameter	*Description*	*Value*
β	Discount factor	0.995
κ	Slope of Phillips curve	0.025
σ	Inverse elasticity of substitution	2
σ_ε	Standard deviation natural rate innovation	1.32
σ_u	Standard deviation of cost-push innovation	0.10
ρ_ε	Serial correlation of natural rate	0.85
ρ_u	Serial correlation of cost-push	0
λ	Weight on output stabilization	0.25
π^*	Steady-state inflation (annualized)	2
ρ_i^n	Value of natural rate at time 1	–0.5
T	Quarters to reach terminal natural rate	24
$\overline{\rho}$	Terminal natural rate (annualized)	1.75
δ	Backward-looking IS curve coefficient	0.75
ξ	Backward-looking Phillips curve coefficient	0.95
x_0	Initial condition for the output gap	–1.5
π_0	Initial condition for inflation	1.3
φ	Taylor rule coefficient on inflation	1.5
γ	Taylor rule coefficient on output gap	0.5

Source: Authors' calculations.

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

coefficient is set midway between the values in Adam and Billi (2007) and Cúrdia and others (2015). We set the standard deviation of the cost-push shock σ_u close to the value used in Adam and Billi (2007). Assuming serial correlation or a moderately different unconditional standard deviation of the cost-push shock is not very important for our results. Finally, by assuming that the economy is not subject to shocks for $t > T$ and that the long run natural rate $\overline{\rho}$ is a known constant, we have been conservative in our specification of uncertainty.

The Phillips curve slope, elasticity of intertemporal substitution, and discount factor are all set to values common in the New Keynesian literature. For the backward-looking model we set the coefficient on lagged inflation in equation 10 to $\xi = 0.95$, reflecting the fact that inflation has been very persistent in recent years.[24] The coefficient on lagged output in equation 11 is $\delta = 0.75$, in order to generate significant persistence in the output gap. For the backward-looking model we assume an initial inflation rate of 1.3 percent, a recent reading for core PCE inflation, and an

24. Note that it is not clear how to map estimates of the lagged inflation coefficient in the literature to our backward-looking model since these are based on Phillips curves with forward-looking terms.

initial output gap $x_0 = -1.5$ percent, based on a simple calculation using the 2014Q4 unemployment rate (5.7 percent), an estimate of the natural rate of unemployment (5.0 percent), and Okun's law. As indicated by Proposition 3, adding uncertainty about the initial output gap would only strengthen our results.[25]

We measure the quantitative effect of uncertainty on policy by comparing equilibrium outcomes under optimal discretion to a scenario in which we solve for optimal discretion when the central bank observes the current natural rate and cost-push shocks but acts as if there will be no more shocks. Private agents understand this policy but take into account the true nature of uncertainty. Actual outcomes will be inconsistent with the central bank's assumptions, so we call this the "naive" policy. We also compare equilibrium outcomes under optimal discretion to those obtained assuming the central bank follows a reaction function with weights on inflation and the output gap as in John Taylor (1993), and a constant term equal to 3.75 percent corresponding to $\rho + \pi^*$.

RESULTS FOR THE FORWARD-LOOKING MODEL Figure 1 displays representative paths of the nominal interest rate, inflation, and the output gap under optimal discretion, the naive policy, and the Taylor rule, calculated by setting the ex post realized shocks to zero, the modal outcome. Under the modal outcome, the interest rate under the naive policy follows the natural rate exactly. The difference between the interest rate paths indicates the substantial impact uncertainty has on optimal policy; the naive policy is between 50 and 150 basis points above the optimal policy for 2 years. This difference in policy has little impact on the output gap, but under optimal policy the inflation gap is closed much faster. The inflation gap is more negative under the naive policy because the interest rate is higher both initially and in the future since it does not take into account uncertainty about the ZLB.[26] The Taylor rule prescribes rates above both the optimal and naive policies for most of the simulation period, and because agents are forward looking this feeds backward to cause much more negative gaps.[27]

Table 2 summarizes the distribution of outcomes under the three different policies based on simulating 50,000 paths drawn from the calibrated

25. In the online appendix we discuss the implications for our results of different values for the initial gaps, uncertainty, ρ^n, δ, and ξ.

26. One might be surprised that inflation is far below target under the naive policy even though the output gap is near the target. This reflects the fact that we plotted the modal outcome, rather than the mean, and that the distributions of inflation and output gap outcomes are skewed to the left.

27. For some calibrations, the outcomes under the Taylor rule can be so poor that liftoff is delayed and rates are below the optimal policy throughout the simulation period.

Figure 1. Liftoff in the Forward-Looking Model

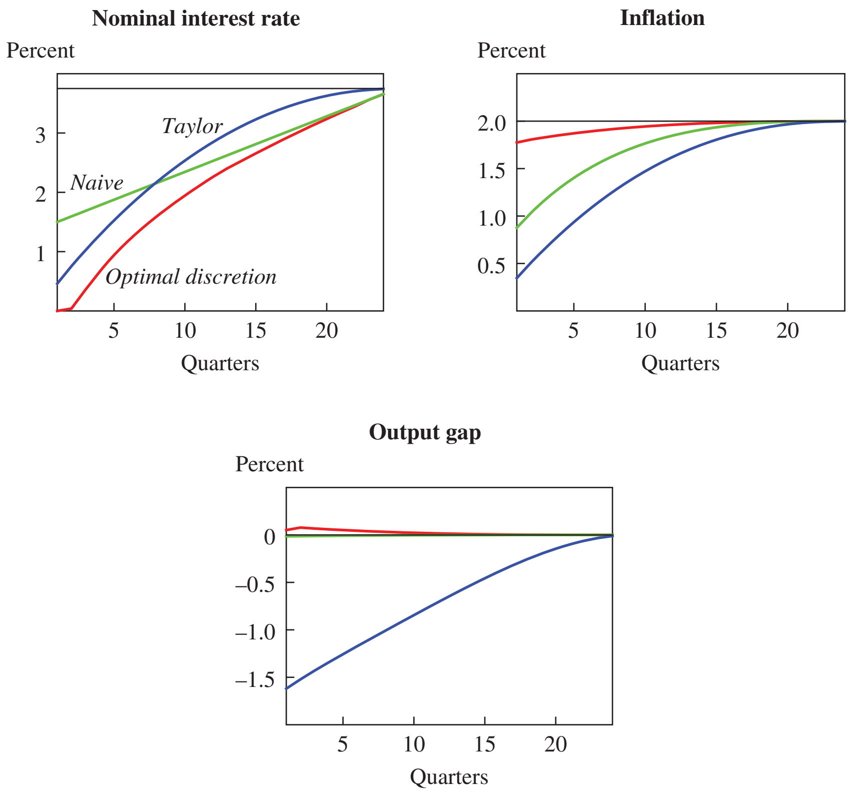

Source: Authors' calculations.

Table 2. Forward-Looking Simulation

Statistic	*Optimal discretion*	*Naive*	*Taylor rule*
Expected loss	0.02	0.06	0.16
Mean time at liftoff	4.11	1.00	1.00
Median time at liftoff	3	1	1
Median π at liftoff	1.81	0.88	0.35
Median x at liftoff	0.08	–1.44	–1.62
75th percentile maximum (π)	2.69	2.42	2.17
25th percentile minimum (x)	–0.72	–1.44	–2.63
Median standard deviation Δi	1.87	1.88	0.97

Source: Authors' calculations.

distributions of the shocks. Optimal discretion implies one-third the loss expected under the naive policy and one-eighth the loss expected under the Taylor rule.[28] One way to interpret these losses is to calculate the per-period reduction in the output gaps and inflation gaps that would make the central banker indifferent between the outcomes under the optimal policy and those under the alternatives. Both gaps would have to be 43 percent and 65 percent smaller under the naive policy and the Taylor rule, respectively, to achieve this indifference. Under optimal discretion, the median liftoff (defined as the nominal interest rate exceeding 25 basis points) is delayed by 2 quarters compared to the other policies; the mean liftoff is delayed by more than 3 quarters, reflecting skewness in the outcomes. At the time of liftoff, inflation and output are much closer to the target under optimal discretion compared to the two alternative policies.

When comparing policies it is also important to assess how well each balances the risks of bad outcomes. We do this by comparing the 75th percentile across simulations of the maximum inflation gap and the 25th percentile of the lowest output gap over the first 6 years. Under optimal policy, the bad output outcomes are much lower than under either alternative policy. The bad inflation outcomes do not seem particularly high under any of the policies.

The statistic in the bottom row is the median standard deviation of changes in the nominal interest rate. By comparing interest rate volatility under the Taylor rule in our model with that implied by the same Taylor rule in the data, we can determine whether the uncertainty underlying our results is reasonable. If the volatility were much higher in our simulations we would conclude that it is unreasonably large. In fact, the 0.97 standard deviation in our Taylor rule simulations is only a little larger than the 0.88 standard deviation we find in our data.[29] Interest rates are more volatile under both the optimal and naive policies because they respond to all fundamental shocks rather than to inflation and output alone.[30]

RESULTS FOR THE BACKWARD-LOOKING MODEL Figure 2 is the analog of figure 1 for the backward-looking model. The dynamics of return to target are quite different from those in the forward-looking model, but the key qualitative results are the same. As in the forward-looking model, optimal

28. The suboptimality of the Taylor rule does not hold by definition, because it provides commitment, which may lead to more favorable outcomes.

29. The online appendix describes how we calculate the interest rate implied by the Taylor rule with our data.

30. We thank Johannes Wieland for suggesting that we assess the volatility of the nominal interest rate.

Figure 2. Liftoff in the Backward-Looking Model

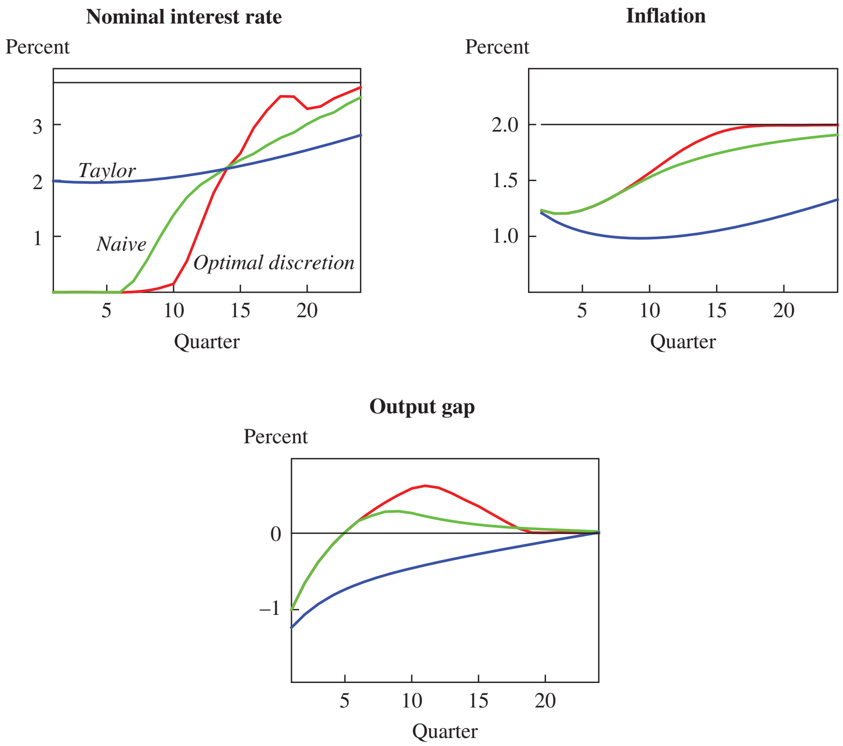

Source: Authors' calculations.

policy is substantially looser than both the naive policy and the Taylor rule. Here the optimal policy prescribes much more delay in lifting off from the ZLB. Delay now occurs under the naive policy because it is optimal to stimulate output strongly in order to return inflation to target, but this delay is shorter than under the optimal policy. The optimal policy also has a sharper liftoff than the naive policy. However, the increases under optimal policy are equivalent to just 25 basis points at each FOMC meeting, the same as the "measured pace" followed during the Fed tightening over 2004–06. Qualitatively, the differences in the output and inflation outcomes across the three policies are similar to those in the forward-looking model as well. Taking into account uncertainty about the ZLB leads the optimal policy to return inflation to target faster than the naive policy, and it achieves this by allowing the output gap to overshoot more in order to build a buffer against the possibility of bad shocks in the future.

Table 3. Backward-Looking Simulation

Statistic	*Optimal discretion*	*Naive*	*Taylor rule*
Expected loss	0.27	0.28	0.60
Mean time at liftoff	12.5	10.3	1.00
Median time at liftoff	10	7	1
Median π at liftoff	2.00	1.81	1.21
Median x at liftoff	0.32	0.00	−1.27
75th percentile max(π)	3.02	2.83	2.81
25th percentile min(x)	−1.65	−1.70	−1.54
Median standard deviation Δi	2.96	3.10	0.54

Source: Authors' calculations.

Table 3 is constructed analogously to table 2. It shows that optimal policy provides only a marginal improvement over the naive policy in terms of expected losses, due to the offsetting effects of the inflation and output gaps. The median gaps are roughly closed at liftoff under both the optimal and naive policies, but they are quite large under the Taylor rule. The bad outcomes are similar across the three scenarios. Finally, note that the volatility of the interest rate under the Taylor rule is lower here compared to the data and the forward-looking model, so the underlying uncertainty is not excessive.

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 liftoff. Figure 3 depicts a particular simulation where there is a large positive cost-push shock before the liftoff under the optimal policy shown in figure 2. The shock triggers earlier liftoff under the optimal policy so that the inflation response is mild. The implication is that staying at zero longer under the optimal policy does not impair the ability of the central bank to respond to future contingencies. However, it does have to be prepared to raise rates promptly. We obtain similar results with the forward-looking model.

II. Historical Precedents for Risk Management

The previous section demonstrates that the ZLB justifies a risk management approach to monetary policy. One may question whether following such an approach would be a departure from past FOMC behavior. Clearly, concerns about the ZLB are a relatively recent phenomenon. Nevertheless there are many reasons why a risk management approach can be justified when away from the ZLB, and we begin this section by reviewing them. We then demonstrate that the Federal Reserve has used risk management to justify its policy decisions over the period 1987–2008.

Figure 3. Large Cost-Push Shock in the Backward-Looking Model

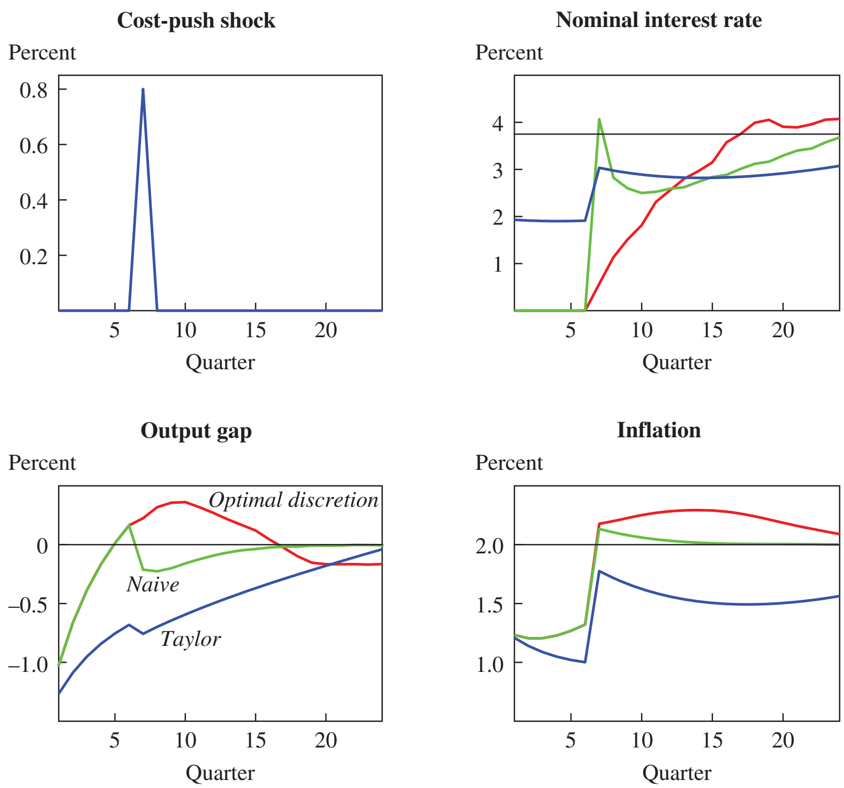

Source: Authors' calculations.

The FOMC minutes and other Federal Reserve communications reveal a number of episodes when uncertainty or insurance were used to justify the Fed's policy decisions. Sometimes the FOMC indicated that it had a wait-and-see approach to taking further actions or muted a funds rate move due to its uncertainty over the course of the economy or the extent to which early policy moves had yet shown through to economic activity and inflation. At other times the FOMC said its policy stance was taken in part as insurance against undesirable outcomes; during these times, the FOMC often noted that the potential costs of a policy over-reaction likely were modest compared to the scenario it was insuring against.

Two episodes are particularly revealing. The first is the hesitancy of the FOMC to raise rates in 1997 and 1998 to counter inflationary threats because of uncertainty generated by the Asian financial crisis and the 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, in late 2001 the FOMC's aggressive actions also seemed to be influenced by attention to the risks associated with the ZLB on interest rates.

While the historical record is replete with references suggesting that the policy stance was influenced by uncertainty or insurance motives, this does not establish that risk management actually had a material impact on policy. Therefore, we conclude this section by quantifying these references into variables that we then use in section III to assess the importance of risk management for actual policy decisions.

II.A. Rationales for Risk Management Away from the ZLB

Policymakers have long emphasized the importance of uncertainty in their decision making. As Alan Greenspan (2004) put it: "The 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."[31] This sentiment seems at odds with linear-quadratic models in which optimal policy involves adjusting the interest rate in response to only the mean of the distribution of shocks away from the ZLB. What kinds of factors cause departures from such conditions and justify the risk management approach?

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 by Woodford (2003) as being a local approximation to consumer welfare. However, it might not be a good approximation when large shocks drive the economy far from the underlying trend; alternatively, it might simply be an inadequate approximation of FOMC behavior. Examples of models with asymmetric loss functions include those described by Paolo Surico (2007), Lutz Kilian and Simone Manganelli (2008), and Juan J. Dolado, P. Ramón María-Dolores, and Francisco Ruge-Murcia (2004).[32] The model studied by the last authors implies that the optimal policy rule can involve nonlinear output gap and inflation terms if policymakers are

31. "Risk and Uncertainty in Monetary Policy," p. 36 (see note 2).
32. For an early contribution of the effects of asymmetric loss functions on stabilization policy see Friedman (1975).

less averse to allowing output to run above potential than below it. The relevance of higher moments in the distribution of shocks for optimal policy is an obvious by-product of these nonlinearities.

Nonlinearities in economic dynamics are another natural motivation. For example, suppose recessions are episodes when self-reinforcing dynamics amplify the effects of downside shocks. This could be modeled as a dependence of current output on lagged output, as in our backward-looking model, but with such dependence being concave rather than linear. Intuitively, negative shocks have a more dramatic effect on reducing future output than positive shocks have on increasing it, so greater uncertainty leads to looser optimal policy to guard against the more detrimental outcomes. Alternatively, suppose the Phillips curve is convex, perhaps owing to downward nominal wage rigidities that become more germane with low inflation. Here, a positive shock to the output gap leads to a significant increase of inflation above target while a negative shock leads to a much smaller decline in inflation. The larger the spread of these shocks, the greater the odds of experiencing a bad inflation outcome. Optimal policy guards against this, leading to a tightening bias.[33] The risk management approach also appears 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 uncertainty over the effects of policy should lead to caution and smaller policy responses to deviations from target. In contrast, the robust control analysis of Lars Hansen and Thomas Sargent (2008) has been interpreted to mean that uncertainty over model mis-specification should generate aggressive policy actions. As explained by Gadi Barlevy (2011), both the attenuation and aggressiveness results depend on the specifics of the underlying environment. Nonetheless, these analyses still often indicate that higher moments of the distribution of shocks can influence the setting of optimal policy.

II.B. 1997–98

The year 1997 was a good one for the U.S. economy: real GDP increased 3¾ percent (the March 1998 third estimate), the unemployment rate fell to 4.7 percent, and core CPI inflation was 2¼ percent. With solid growth and

33. 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).

tight labor markets, 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.[34]

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 funds rate 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 growth to moderate to a more sustainable pace as well as the fact that actual inflation had remained contained despite tight labor market conditions. Based on the funds rate remaining at 5.5 percent, the Board of Governors' staff forecast in the August 1998 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. Additionally, however, on several 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 1997 FOMC minutes is a revealing example:

> 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 FOMC an opportunity to gauge the momentum of the expansion and the related degree of pressure on resources and prices.[35]

Furthermore, the FOMC did not regard "waiting and seeing" as having a high cost. They thought any increase in inflation would be slow and that,

34. Available at http://www.federalreserve.gov/boarddocs/hh/1998/february/ReportSection1.htm

35. Available at http://www.federalreserve.gov/fomc/minutes/19970701.htm

if needed, a limited tightening would be sufficient to rein 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 FOMC's decision to leave policy on hold.

Of course, the potential fallout for the U.S. economy of the Asian financial crisis 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 inflationary pressures in check but would 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.[36]

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 United States. The FOMC was concerned about the direct implications of these developments for U.S. financial markets, already evident in the data, as well as their implications for the real economy, which were still just a prediction. The staff forecast

36. "Coming Budgetary Challenges," Testimony of Chairman Alan Greenspan before the Committee on the Budget, U.S. House of Representatives, March 4, 1998. Available at http://www.federalreserve.gov/boarddocs/testimony/1998/19980304.htm

prepared for the September FOMC meeting reduced the projection for growth in 1999 by about ½ percentage point to 1¼ percent, predicated on a 75 basis-point reduction in the funds rate spread out over three quarters. Such a forecast was not a disaster—indeed, at 5.2 percent the unemployment rate projected for the end of 1999 was still below the staff's estimate of its natural rate. Nonetheless, the FOMC moved much faster than the staff assumed it would, lowering rates 25 basis points at its September and November meetings as well as making an inter-meeting rate 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 of that year:

> 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.[37]

While the references to insurance are clear, a case also can be made that these policy moves were undertaken largely to realign the misses in the expected paths for growth and inflation from the FOMC's policy goals. At that time, the prescriptions to address the risks to their 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. As the above quote from Chairman Greenspan's February 1998 testimony indicated, in early 1998 the FOMC thought that a 5½ percent funds rate kept these risks in balance. Subsequently, as the odds of economic weakness increased, the FOMC cut rates to bring the risks to the two goals back into balance. As Chairman Greenspan said in his February 1999 Humphrey-Hawkins 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. . . .

37. Minutes of the Federal Open Market Committee, September 29, 1998. Available at http://www.federalreserve.gov/fomc/minutes/19980929.htm

> These actions were taken to rebalance the risks to the outlook, and, in the event, the markets have recovered appreciably.[38]

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 III, which seeks to disentangle the normal response of policy to expected outcomes from uncertainty and other related factors that may have influenced the policy decision.

II.C. 2000–01

In the end, the economy weathered the fallout from the Russian default well. The strength of the U.S. economy and underlying inflationary pressures led the FOMC to execute a series of rate hikes that brought the funds rate up to 6.5 percent by May of 2000. At the time of the June 2000 FOMC meeting, the unemployment rate stood at 4.1 percent and core PCE inflation, which the FOMC was now using as its main measure of consumer price inflation, was running at about 1¾ percent, up from 1½ percent in 1999. The staff forecast that growth would moderate to a rate near or a little below potential, the unemployment would remain near its current level, and inflation would rise to 2.3 percent in 2001—and this forecast was predicated on another 75 basis points tightening. 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 would play out. First, the incoming data and anecdotal reports from committee members' business contacts pointed to a slowdown in growth, although how much it was slowing was unclear. Second, with rates having risen substantially over the previous 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 funds rate and the slowdown in growth that appeared in train, the FOMC 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.

38. "The Federal Reserve's Semiannual Report on Monetary Policy," Testimony of Chairman Alan Greenspan before the Committee on Banking, Housing, and Urban Affairs, U.S. Senate, February 23, 1999. Available at http://www.federalreserve.gov/boarddocs/hh/1999/february/testimony.htm

Accordingly, the FOMC thought it better to put a rate increase on hold and see how the economy evolved. The June 2000 FOMC minutes contain a good deal of commentary supporting this interpretation:[39]

> 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.[40]

In the second half of 2000 it became increasingly evident that growth had slowed to a pace somewhat below trend and inflation was moving up at a slower pace than the staff had projected in June. The FOMC's response was to hold the funds rate at 6.5 percent through the end of 2000. But the data around the turn of the year proved to be weaker than anticipated. In a conference call on January 3, 2001, the FOMC cut the funds rate to 6 percent, and then at its end-of-month meeting it lowered the rate again, to 5½ percent.[41]

In justifying the aggressive ease, the minutes 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

39. The FOMC had already invoked such arguments earlier in this cycle. As noted in the July 2000 Monetary Policy Report: "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." (Monetary Policy Report forwarded to Congress on July 20, 2000, available at http://www.federalreserve.gov/boarddocs/hh/2000/July/ReportSection1.htm)

40. Minutes of the Federal Open Market Committee, June 27–28, 2000. http://www.federalreserve.gov/fomc/minutes/20010131.htm

41. At that meeting the Federal Reserve 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 funds rate at 5.75 percent. Core PCE inflation was projected to rise modestly to a little under 2.0 percent.

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.[42]

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 FOMC to make a large move. The “help guard against cumulative weakness” and “front-loaded” language could be read as the FOMC taking out some additional insurance against the possibility that the weakening activity would snowball into a recession. This could have reflected a concern about the kinds of nonlinear output dynamics or perhaps non-quadratic losses associated with a large recession that we discussed in section II.A.

The FOMC steadily brought the funds rate down further over the course of 2001, against a backdrop of weakening activity, and the economy seemed to be skirting a recession. Then the tragic events of September 11 occurred. There was, of course, huge uncertainty 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 1½ percent at an annual rate and rise at just a 1¼ 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 1½ percent. These forecasts were predicated on the funds rate remaining flat at 2¼ percent.

However, in the aftermath of the terrorist attacks the FOMC 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 funds rate already low. In response, the FOMC again acted aggressively, cutting the funds rate 50 basis points in a conference call on September 17 and again at their regular meetings in October and November. The November 2001 FOMC meeting minutes note:

. . . 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

42. Minutes of the Federal Open Market Committee, January 30–31, 2001. Available at http://www.federalreserve.gov/fomc/minutes/20010131.htm

> 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.[43]

This passage suggests that the large rate cuts were not only aimed at preventing the economy from falling into a serious recession with deflationary consequences, but that the FOMC 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 being constrained by the ZLB, precisely the policy scenario and optimal policy prescription described in section I.

II.D. Quantifying References to Uncertainty and Insurance in FOMC Minutes

We have shown that Federal Reserve communications contain many references suggesting that uncertainty or insurance motives influenced the stance of policy. But the question remains: Has risk management had a material impact on policy? We now show how we quantified these references into variables that can be used to assess the importance of risk management for actual policy decisions.

In the spirit of the narrative approach pioneered by Christina Romer and David Romer (1989), we built judgmental indicators based on our reading of the FOMC minutes covering the period from the beginning of Greenspan's chairmanship in 1987 to 2008. We concentrated on the paragraphs that describe the FOMC's rationale for its policy decision, reading these passages for references to when uncertainty or insurance considerations appeared closely linked to the FOMC's decision. Other portions of the minutes were excluded from our analysis in order to better isolate arguments that directly influenced the policy decision from more general discussions of unusual data or forecast uncertainty.

43. Minutes of the Federal Open Market Committee, November 6, 2001. A transcript is available at https://www.andrew.cmu.edu/course/88-301/monetarism/minutes-0111.pdf

We constructed two separate judgmental variables, one for uncertainty (*hUnc*) and one for insurance (*hIns*), where "h" stands for "human-coded." The uncertainty variable was coded to plus (minus) one if we judged that the FOMC appealed to uncertainty to position the funds rate higher (lower) than it otherwise would be based on the staff forecast alone. If uncertainty did not appear to be an important factor influencing the policy decision, we coded the indicator as zero. We coded the insurance variable similarly by identifying when the minutes cited insurance against some adverse outcome as an important consideration in the stance of policy.[44]

As an example of our coding, consider the June 2000 meeting discussed above when the FOMC decided to wait to assess future developments before taking further policy action. The commentary below highlights the role of uncertainty in this decision (our 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 ½ 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.*[45]

We coded this meeting as a minus one for *hUnc*—rates were lower because uncertainty over the economic outlook and the effects of past policy moves appear to have been important factors in the FOMC's decision not to raise rates. Similarly, the January and November 2001 quotes cited above led us to code *hIns* as a minus one for those meetings, since, as we noted in the narrative, the FOMC appeared to be making aggressive rate moves in part to insure against downside risks to the baseline scenario.

We did not code all mentions of uncertainty or insurance as a plus or minus one. For example, the March 1998 minutes referred to uncertainties over the economic outlook and said that the FOMC could wait for further developments before tightening to counter potential inflation developments.

44. A value of plus (minus) one for either variable could reflect the FOMC raising (lowering) rates by more (less) than they would have if they ignored uncertainty or insurance *or* a decision to keep the funds rate at its current level when a forecast-only call would have been to lower (raise) rates.

45. See note 40.

However, at that time the FOMC was not obviously in the midst of a tightening cycle; the baseline forecast seemed consistent with the funds rate setting at the time; and the commentary over the need to tighten was in reference to an indefinite point in the future. So, in our judgment, uncertainty did not appear to be a very important factor holding back a rate increase at that meeting, and we coded it as a zero.[46]

Of course, this coding of the minutes is inherently subjective, and there is no definitive way to judge the accuracy of the decisions we made. Consequently 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 variables which measure the percentage of sentences containing words related to uncertainty or insurance in conjunction with references to economic activity, inflation, or both.[47] The measures for uncertainty and insurance are denoted *mUnc* and *mIns*, where "m" indicates these variables are "machine-coded." Figures 4 and 5 show plots of our minutes-based uncertainty and insurance variables.

Non-zero values of the human-coded variables are indicated by dots and the bars indicate the machine-coded sentence counts. The uncertainty indicator *hUnc* "turns on" in 31 out of the 128 meetings between 1993 and 2008. Indications that insurance was an actor in shading policy are not as common, but still show up 14 times in *hIns*. Most of the time—24 for uncertainty and 11 for insurance—we judged that rates were set lower than they otherwise would have been to account for these factors.

The *hUnc* and *hIns* codings are not always reflected in the sentence counts. There are also meetings where the sentence counts are positive but we did not judge them to indicate that rates were set differently than they normally would have been. For example, in March 2007 *hUnc* is coded zero for uncertainty whereas *mUnc* finds uncertainty referenced in nearly one-third of the sentences in the policy section of the minutes. Inspection of the minutes indicates that the FOMC was uncertain over both the degree to which the economy was weakening and whether their expectation of a decline in inflation, which was running uncomfortably high at the time,

46. From the minutes: "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 . . . (o)n 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."

47. The appendix describes our coding algorithm in more detail.

Figure 4. Minutes-Based Uncertainty Variables

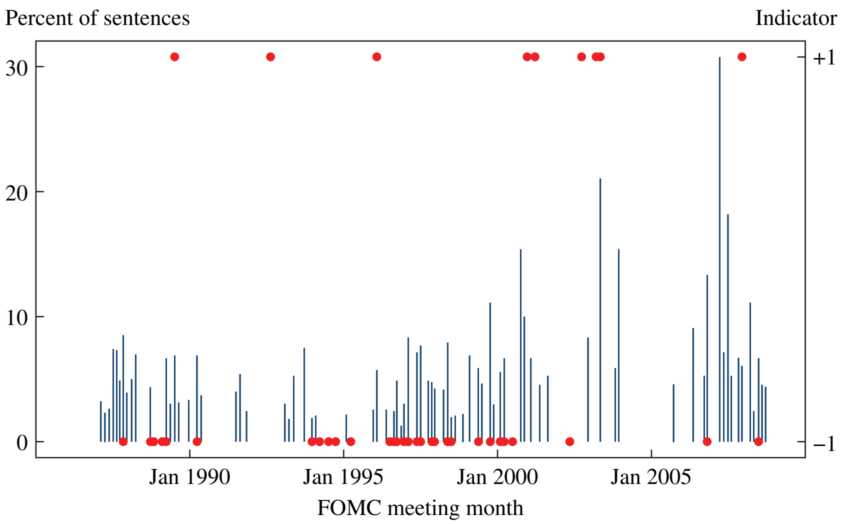

Source: Authors' calculations based on FOMC minutes; see text.

Figure 5. Minutes-Based Insurance Variables

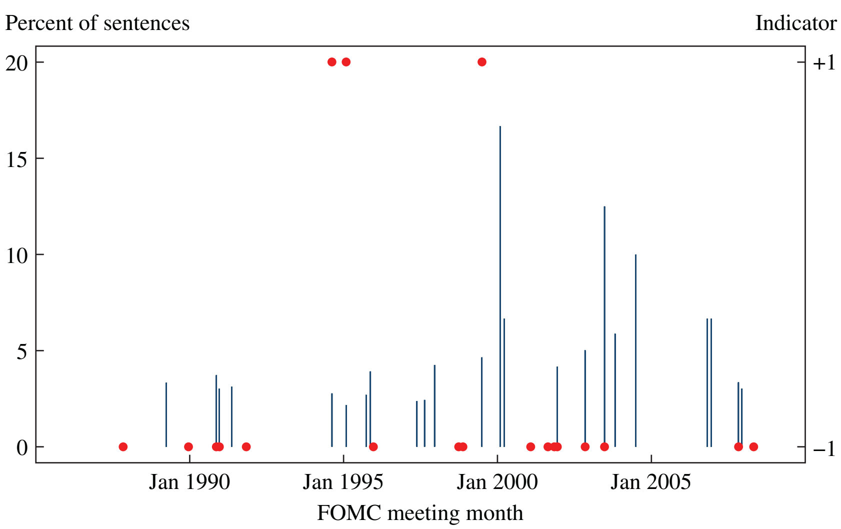

Source: Authors' calculations based on FOMC minutes; see text.

actually would materialize. In the end, they did not adjust current policy in response to these conflicting uncertainties. Hence we coded *hUnc* to zero in this case.

Note that we did not attempt to measure a variable for risk management per se. The minutes often contain discussions of policies aimed at addressing risks to attaining the FOMC's goals. However, many times this commentary appears to surround policy adjustments aimed instead at balancing (possibly conflicting) risks to the outlook for output and inflation, not unlike the response to changes in economic conditions prescribed by the canonical framework for studying optimal policy under discretion. Such risk balancing was discussed in our narrative of the 1997–98 period.[48]

III. Econometric Evidence of Risk Management

So far we have uncovered clear evidence that risk management considerations have been a pervasive feature of Federal Reserve communications. But it is not clear at this stage whether risk management has had a material impact on the FOMC's policy decisions. If it has, then calling for a risk management approach in the current policy environment would be consistent with a well established approach to monetary policy. In this section we describe econometric evidence suggesting that risk management has had a material impact on the FOMC's funds rate choices in the pre-ZLB era.

We estimate monetary policy reaction functions of the kind studied by Clarida, Galí, and Gertler (2000) and many others. These have the funds rate set as a linear function of output gap and inflation forecasts; there is no role for risk management unless risk feeds directly into the point forecasts. To quantify the role of risk beyond such a direct influence we add variables that proxy for risk to the reaction function.[49]

48. Indeed, for much of our sample period, the FOMC discussed risks about the future evolution of output or inflation 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: "An asymmetric directive was consistent with their view that the risks clearly were in the direction of excessive demand pressures." Since the FOMC 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. But it does not appear to be a reference that variance or skewness in the distribution of possible inflation outcomes should dictate some non-standard policy response.

49. There is a large literature that examines nonlinearities in policy reaction functions (see Gnabo and Moccero [2015], 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.

III.A. Empirical Strategy

Let R_t^* denote the *notional target* for the funds rate in period *t*. We assume the FOMC sets this target according to

$$R_t^* = R^* + \beta\left(E_t\left[\pi_{t,k}\right] - \pi^*\right) + \gamma E_t\left[x_{t,q}\right] + \mu s_t, \tag{13}$$

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 denotes expectations conditional on information available to the FOMC at date t. The coefficients β, γ, and μ are fixed over time. R^* is the desired nominal rate when inflation is at target, the output gap is closed, and risk does not influence policy other than through the forecast, $\mu = 0$. If the average output and inflation gaps are both zero and the FOMC acts as if the natural rate is constant and out of its control, then $R^* = r^* + \pi^*$, where r^* is the real natural rate of interest.[50]

We make two more assumptions to arrive at our estimation equation. First, the FOMC has a preference for interest rate smoothing and so does not choose to hit its notional target instantaneously, and as a practical matter it is necessary to include lags of the funds rate to fit the data. Second, the FOMC does not have perfect control over interest rates, which gives rise to an error term, υ_t. These assumptions lead to the following specification for the actual funds rate, R_t:

$$R_t = (1 - A(1))R_t^* + A(L)R_{t-1} + \upsilon_t, \tag{14}$$

where $A(L) = \sum_{j=0}^{N-1} a_{j+1}L^j$ is a polynomial in the lag operator L with N denoting the number of funds rate lags. The error term υ_t is assumed to be mean zero and serially independent. Combining equations 13 and 14 yields our estimation equation:

$$R_t = b_0 + b_1 E_t\left[\pi_{t,k}\right] + b_2 E_t\left[x_{t,q}\right] + A(L)R_{t-1} + b_3 s_t + \upsilon_t, \tag{15}$$

where b_i, $i = 0, 1, 2, 3$ are simple functions of $A(1)$, β, γ, μ, r^* and π^*.[51]

50. There is no presumption that (equation 13) 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.

51. We make no attempt to address 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.

We use the publicly available Federal Reserve 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 $E_t\ [\pi_{t,k}]$ and $E_t\ [x_{t,q}]$ with $k = q = 3$.[52] These forecasts are available for every FOMC meeting. We estimate equation 15 both meeting-by-meeting and quarter-by-quarter.[53] When we estimate it at the quarterly frequency we use staff forecasts corresponding to FOMC meetings closest to the middle of each quarter.[54] We measure R_t at the meeting frequency using the funds rate target announced (or estimated) at the end of the day of a meeting, and we measure it at the quarterly frequency using the average effective funds rate over the 30 trading days following the meeting closest to the middle of the quarter. Provided the error term υ_t is serially uncorrelated and is orthogonal to the forecasts and the risk proxies, we can obtain consistent estimates of β, γ, and μ by estimating equation 15 by ordinary least squares. We keep N sufficiently large to ensure that υ_t is serially uncorrelated.

To quantify the role of risk we study the magnitude and statistical significance of estimates of μ in equation 13. An insignificant estimate of μ cannot be interpreted as evidence against a role for risk management, because risk might operate by influencing point forecasts as in our forward-looking model. We also could find no effect because risk might tilt policy in opposite directions depending on the circumstances. With the exception of our human-coded FOMC-based variables, none of our risk proxies accounts for the fact that perceived risks to the forecast might have different effects on policy depending on the nature of the risk and the state of the economy. For example, an increase in uncertainty about the inflation outlook should lead to tight policy if this increase occurs during a period of heightened concerns about rising inflation, but to looser policy if concerns are over unwanted disinflation. As such, estimates of the effect of any given proxy will at best reflect the nature of the risk and the circumstances in which it has arisen that have predominated over the sample period.

Finally, we do not allow for the coefficients on the forecasts to depend on our risk proxies as is suggested by the work of Brainard (1967) and others. However, we show in the online appendix that if these forecast coefficients are linear functions of risk, then the null hypothesis that a given

52. The online appendix describes our data in more detail.

53. We assume meetings are equally spaced even though this is not true in practice. We account for this discrepancy when we calculate standard errors by allowing for heteroskedasticity.

54. Gnabo and Moccero (2015) also estimate quarterly reaction functions using board staff forecasts.

proxy's coefficient is zero in our now mis-specified model encompasses the null that the forecast coefficients are invariant to risk as measured by that proxy.

III.B. Proxies for Risk Management

In addition to our human- and machine-coded FOMC-based variables we consider several proxies for risk management that do not rely on interpreting the FOMC minutes. Two of these variables are constructed using the Federal Reserve Board staff's forecast, which is seen by the FOMC at its regular meetings, and we study them using our meeting frequency reaction functions. The remaining variables are measured at the quarterly frequency and can be divided into two groups based on whether they primarily reflect variance or skewness in the forecast.

The two additional FOMC-based proxies involve revisions to the Federal Reserve Board staff's forecasts for the output gap (frGap) and core CPI inflation (frInf). The revisions correspond to changes between meeting m and $m - 1$ in the forecasts over the same one-year period that starts in the quarter of meeting $m - 1$. A big change in the forecast is usually triggered by unusual events that may be difficult to interpret and hence generate uncertainty about the forecast. If the FOMC were only worried about these events in making its point forecast, then the post-shock forecasts of the output gap or inflation would be sufficient to describe the policy setting. However, if uncertainty has a separate effect on policy the forecast revisions might enter significantly.

Three of the quarterly proxies exploit financial market data: VXO, SPD, and JLN. VXO is the Chicago Board Options Exchange's measure of market participants' expectations of volatility in the S&P 500 stock index over the next 30 days. Since the S&P 500 reflects earnings expectations, VXO should, at least in part, measure market participants' uncertainty about the economic outlook.[55] SPD is 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 private-sector default risk plus other factors that may indicate downside risks to economic growth.[56] JLN is Kyle Jurado, Sydney Ludvigson, and

55. Using a VAR framework Bekaert, Hoerova, and Lo Duca (2013) find weak evidence that positive innovations to VXO lead to looser policy. Gnabo and Moccero (2015) find that policy responds more aggressively to economic conditions and is less inertial in periods of high uncertainty as measured by VXO.

56. Alcidi, Flamini, and Fracasso (2011), Castelnuovo (2003), and Gerlach-Kristen (2004) consider reaction functions including SPD.

Serena Ng's (2015) measure of the common variation in the one-year-ahead unforecastable components of a large number of activity, inflation, and financial indicators. Given its basis in measuring uncertainty about macroeconomic forecasts, JLN is a natural risk proxy to consider. But, unlike VXO and SPD, it does not measure real-time uncertainty, and similar to these two measures it confounds macroeconomic and financial uncertainty.

The remaining proxies are based on the Survey of Professional Forecasters (SPF) which surveys forecasters about their point forecasts of GDP growth and GDP deflator inflation and their probability distributions for these forecasts. We use both kinds of information to construct measures of variance and skewness in the economic outlook one year ahead.[57] Variance is measured using the median among forecasters of the standard deviations calculated from each individual's probability distribution (vGDP and vInf) and the interquartile range of point forecasts across individuals (DvGDP and DvInf.)[58] Skewness is measured using the median of the individual forecasters' mean minus mode (sGDP and sInf) and the difference between the mean and the mode of the cross-forecaster distribution of point forecasts (DsGDP and DsInf). Consequently, a positive (negative) value for one of these proxies represents upside (downside) risk to the modal forecast.

The principal advantage of these proxies is that they are real-time measures of perceived risks in the forecast. The main drawback of the measures based on survey respondents' forecast distributions is that the bins they are asked to put probability mass on are relatively wide, so statistics based on them may contain substantial measurement error. The proxies based on the cross-section of forecasts are properly thought of as measuring forecaster disagreement rather than variance or skewness in the outlook per se. However, there is a large literature that uses forecaster disagreement as a proxy for perceived risk.[59]

All estimates are based on samples that end in 2008 to avoid the ZLB period but begin at different dates to address idiosyncratic features of the data. The benchmark start date is determined by the onset of Alan Greenspan's tenure as chairman of the FOMC in 1987, but later dates are used in several cases. The sample for the FOMC-based indicators starts in 1993

57. The forecast distributions are for growth and inflation in the current and following year. We use D'Amico and Orphanides' (2014) procedure to translate these into distributions of four-quarter-ahead forecasts.

58. Gnabo and Moccero (2014) find statistically insignificant effects of DvInf on monetary policy.

59. As discussed by Baker, Bloom, and Davis (2015) there is no consensus on how good a proxy it is. Note that we do not study Baker, Bloom, and Davis's (2015) measure of uncertainty since it confounds uncertainty about monetary policy and the economic outlook.

because inter-meeting changes in the target funds rate were much more common prior to that year than afterwards; the FOMC often voted on a bias to future policy moves and the chairman subsequently acted at his discretion. We cannot use inter-meeting moves because we lack contemporaneous staff forecasts. Furthermore, the change in the frequency of inter-meeting moves raises the spectre of instability in the reaction function.[60] The pre-1993 inter-meeting moves are less of a concern for our quarterly models, because in these specifications the funds rate is not as closely tied to any particular meeting. So we chose to include these data points to maximize the number of observations, except when considering the proxies based on individuals' forecast distributions from the SPF. In the latter cases, the first observation is 1992Q1, to coincide with a discrete change in SPF methodology.[61]

Tables 4 and 5 display summary statistics for Federal Reserve Board staff forecasts of inflation and the output gap and the various proxies for risk management at the meeting and quarterly frequencies. What is most worth noting in these tables is that no risk proxy displays a particularly large positive or negative correlation with either the output gap or inflation forecast. This suggests that our proxies contain information that is not already incorporated into these forecasts. Nevertheless, some variables have moderately large correlations in absolute value, so the forecasts do somewhat reflect underlying risks to the outlook. Interestingly, skewness in forecasters' GDP forecasts (GDP) is negatively correlated with the outlook for activity.

Tables 6 and 7 display cross-correlations of the FOMC-based and quarterly proxies, respectively. As suggested by figures 4 and 5, the human- and machine-coded FOMC variables for uncertainty and insurance are essentially uncorrelated. These variables also appear unrelated to the forecast revision variables. However, several correlations among the quarterly proxies are worth noting. Forecaster variance and disagreement about the GDP growth outlook (vGDP and DvGDP) are both positively correlated

60. Between 1990 and 1992, only 4 of the 18 changes in the funds rate target occurred at an FOMC meeting. In contrast, between 1993 and 2008, 54 of the 61 changes in the funds rate target occurred at FOMC meetings. Ignoring inter-meeting moves causes specification problems if interest rate smoothing is a function not only of time but also of the number of policy moves. Indeed, when we estimated our meeting frequency models starting in 1987, our point estimates were (statistically) similar, but even with 5 funds rate lags substantial serial correlation remained in the residuals.

61. In 1992 the SPF narrowed the bins it used 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. Summary Statistics for the FOMC-Based Risk Proxies

						Correlation with forecast of	
Variable	*Obs.*	*Mean*	*Standard deviation*	*Minimum*	*Maximum*	*Inflation*	*Output gap*
Inflation forecast	128	2.45	0.45	1.30	3.53	1.00	0.21
Output gap forecast	128	–0.14	1.58	–4.85	3.08	0.21	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	4.80	0	30.8	–0.06	0.14
mIns	128	0.83	2.45	0	16.7	–0.10	0.08
frInf	128	–0.01	0.18	–0.63	0.63	0.23	0.01
frGap	128	–0.01	0.41	–2.00	0.77	0.24	0.29

Source: Authors' calculations, based on Philadelphia Fed Greenbook data sets and FOMC minutes; see text.

Table 5. Summary Statistics for Quarterly Risk Proxies

						Correlation with forecasts of	
Variable	*Obser-vations*	*Mean*	*Standard deviation*	*Minimum*	*Maximum*	*Inflation*	*Output Gap*
Inflation forecast	86	2.97	1.02	1.33	5.32	1.00	–0.04
Output gap forecast	86	–0.45	1.69	–4.4	3.08	–0.04	1.00
VXO	86	21.0	8.48	10.6	62.1	–0.02	0.04
JLN	86	0.96	0.05	0.89	1.22	–0.06	–0.04
vInf	68	0.74	0.06	0.6	0.90	–0.22	–0.08
vGDP	68	0.9	0.12	0.67	1.30	–0.22	0.22
DvInf	86	0.6	0.18	0.24	1.10	0.25	–0.35
DvGDP	86	0.73	0.27	0.3	1.64	0.37	–0.05
SPD	86	2.11	0.65	1.37	5.60	–0.34	–0.34
sInf	68	0.05	0.08	–0.12	0.30	0.23	–0.12
sGDP	68	–0.10	0.19	–0.54	0.47	–0.10	–0.48
DsInf	86	0.06	0.20	–0.5	0.51	0.01	–0.23
DsGDP	86	0.3	0.27	–0.5	0.90	–0.22	0.21

Source: Authors' calculations, based on Philadelphia Fed Greenbook data sets, Survey of Professional Forecasters, Haver Analytics, and Jurado, Ludvigson, and Ng (2015); see text.

Table 6. Cross-Correlations of FOMC-Based Risk Proxies

Variable	*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	0.25

Source: Authors' calculations, based on Philadelphia Fed Greenbook data sets and FOMC minutes; see text.

Table 7. Cross-Correlations of Quarterly Risk Proxies

Variable	*VXO*	*JLN*	*vInf*	*vGDP*	*DvInf*	*DvGDP*	*SPD*	*sInf*	*sGDP*	*DsInf*
JLN	0.54									
vInf	0.04	0.23								
vGDP	0.40	0.29	0.40							
DvInf	0.15	0.18	0.29	0.03						
DvGDP	0.54	0.38	0.16	0.37	0.33					
SPD	0.73	0.67	0.15	0.32	0.26	0.35				
sInf	–0.27	–0.11	0.29	–0.16	0.08	–0.14	–0.18			
sGDP	0.21	0.22	–0.09	–0.04	0.25	0.16	0.43	–0.15		
DsInf	–0.28	–0.17	0.10	–0.24	0.13	–0.14	–0.08	0.04	0.15	
DsGDP	0.07	–0.08	0.04	0.10	–0.22	0.08	0.02	–0.08	–0.17	–0.17

Source: Authors' calculations, based on Survey of Professional Forecasters; Jurado, Ludvigson, and Ng (2015); and Haver Analytics. See text.

with VXO and SPD, suggesting that the financial variables do reflect some uncertainty about the growth outlook. Also, the relatively high correlation of SPD with sGDP suggests that the former to some extent captures skewness in the growth outlook. The correlation of vGDP with vInf and DvGDP with DvInf are both fairly large, suggesting that uncertainty about inflation and uncertainty about GDP often move together. The correlations of the corresponding forecaster uncertainty and disagreement variables (vGDP with DvGDP and vInf with DvInf) are somewhat large too. Evidently, the amount of disagreement among forecasters is similar to the median amount of uncertainty they see. Finally Jurado, Ludvigson, and Ng's (2015) measure of macroeconomic uncertainty, JLN, is highly correlated with VXO and SPD and to some extent with DvGDP, but much less so with any of the other risk proxies.

III.B. Policy Rule Findings

Table 8 shows our policy rule estimates with and without the various FOMC-based variables. Tables 9 and 10 show estimates with and without the quarterly variance and skewness proxies. Except for the human-coded variables *hUnc* and *hIns*, prior to estimation the risk proxies have been normalized to have mean zero and unit standard deviation, so their coefficients indicate percentage-point responses of the funds rate to standard deviation changes. The tables have the same layout: the first column shows the policy rule excluding any risk proxies, and the other columns show the policy rules after adding the indicated risk proxy. The coefficient associated with a given risk proxy corresponds to an estimate of μ in equation 13. The speed of adjustment to the notional funds rate target $\left(\sum_{j=1}^{N} a_j\right)$ and the coefficients on the forecasts of inflation (β) and the output gap (γ) are similar across specifications and consistent with estimated forecast-based policy rules in the literature.

From table 8 we see that the coefficient on the human coding of uncertainty (*hUnc*) is statistically significant at the 5 percent level, indicating that when uncertainty has shaded the policy decision above or below the forecast-only prescription it has moved the notional target by 40 basis points. With interest rate smoothing the immediate impact is much smaller; the 95 percent confidence interval is 2 to 14 basis points. The machine coding of uncertainty (*mUnc*) is significant at the 10 percent level but the effect is small. The insurance indicators (*hIns* and *mIns*) are not significant, but the point estimate of the *hIns* coefficient is similar to its uncertainty counterpart. The coefficient on the output gap forecast revision variable (frGap) is large and significant, indicating a one-standard-deviation positive surprise

Table 8. FOMC-Based Risk Proxies in Monetary Policy Rules[a]

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*
$\sum_{j=1}^{5} a_j$	.81***	.80***	.81***	.80***	.81***	.84***	.81***
	(.03)	(.03)	(.03)	(.03)	(.03)	(.03)	(.04)
β	1.89***	1.95***	1.86***	1.90***	1.89***	1.75***	1.89***
	(.17)	(.16)	(.17)	(.17)	(.17)	(.22)	(.17)
γ	.85***	.88***	.85***	.83***	.85***	.80***	.85***
	(.05)	(.05)	(.05)	(.05)	(.05)	(.06)	(.06)
hUnc		.40**					
		(.16)					
hIns			.48				
			(.45)				
mUnc				.11*			
				(.06)			
mIns					−.0006		
					(.05)		
frGap						.47**	
						(.19)	
frInf							−.009
							(.14)
LM[b]	.31	.07	.59	.58	.31	.63	.20
Obs.	128	128	128	128	128	128	128

Source: Authors' calculations, based on Philadelphia Fed Greenbook data sets and FOMC minutes; see text.

a. Standard errors are robust to heteroskedasticity. Statistical significance at the ***1, **5, and *10 percent levels.

b. Entries in the "LM" row are *p*-values of Durbin's test for the null hypothesis of no serial correlation in the residuals up to the fifth order.

in the forecast raises the notional target by 47 basis points over and above the impact this surprise has on the forecast itself.[62] In contrast, revisions to the inflation outlook (frInf) do not influence policy beyond their direct effect on the forecast.

Table 9 shows clear evidence that variance in the economic outlook has shaded policy away from the forecast-only prescription. The coefficients on VXO and JLN are both statistically and economically significant, with one-standard-deviation increases lowering the notional target funds rate by 43 and 29 basis points, respectively.[63] Disagreement over the GDP forecast (DvGDP) has a significant coefficient, which is similar to the ones for VXO and JLN, suggesting that the latter variables' correlation with monetary policy reflects uncertainty in the growth outlook. That all these coefficients are negative suggests that higher uncertainty about growth has influenced the FOMC when it was concerned about recessionary dynamics and lowered the funds rate more than prescribed by the forecast alone. The only other significant coefficient in table 9 corresponds to the measure of individual forecasters' views about the uncertainty in their inflation forecasts (vInf). In this case uncertainty shades the policy higher, by about 20 basis points. This suggests that higher uncertainty about the inflation forecast has influenced the FOMC when it was concerned about inflation rising above desired levels and raised rates above levels prescribed by the baseline forecast.

Similarly strong evidence that skewness has mattered for policy decisions is found in table 10. The coefficients are significant on the interest-rate-spread indicator of downside risks to activity (SPD), skewness in the outlook for inflation measured from forecasters' own forecast distributions (sInf), and skewness in the inflation outlook measured across point forecasts (DsInf). An increase in perceived downside risks to activity lowers the funds rate, while an increase in perceived upside risks to inflation raises it. The effects seem large; increases in the skewness proxies change the notional target by −56, 23, and 40 basis points, respectively.

These findings reinforce our findings on the variance proxies and, similarly, seem consistent with our reading of FOMC communications. The

62. The magnitude and significance of this coefficient is largely driven by the sharp decline in the funds rate in 2008 that occurred alongside substantial downward revisions to the output gap forecast.

63. The JLN variable can be expressed as a linear combination of the three uncertainty measures constructed with the underlying activity, inflation, and financial indicators separately. We used Jurado, Ludvigson, and Ng's (2015) replication software to separate out these components, and found that the estimated effects of JLN are driven primarily by the financial indicators.

Table 9. Quarterly Variance Proxies in Monetary Policy Rules[a]

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*
$\sum_{j=1}^{2} a_j$	.69***	.69***	.70***	.70***	.69***	.69***	.69***
	(.03)	(.03)	(.03)	(.04)	(.04)	(.04)	(.03)
β	1.73***	1.73***	1.72***	2.21***	2.13***	1.75***	1.88***
	(.12)	(.11)	(.12)	(.17)	(.16)	(.11)	(.13)
γ	.80***	.84***	.81***	.78***	.77***	.78***	.81***
	(.06)	(.06)	(.06)	(.07)	(.06)	(.07)	(.06)
VXO		−.43***					
		(.11)					
JLN			−.29***				
			(.09)				
vInf				.21**			
				(.10)			
vGDP					.03		
					(.12)		
DvInf						−.09	
						(.13)	
DvGDP							−.38***
							(.13)
LM[b]	.53	.56	.86	.71	.59	.52	.86
Obs.	86	86	86	68	68	86	86

Source: Authors' calculations, based on Philadelphia Fed Greenbook data sets, Survey of Professional Forecasters, Haver Analytics, and Jurado, Ludvigson, and Ng (2015); see text.

a. Standard errors are robust to heteroskedasticity. Statistical significance at the ***1, **5, and *10 percent levels.

b. Entries in the "LM" row are p-values of Durbin's test for the null hypothesis of no serial correlation in the residuals up to the second order.

Table 10. Quarterly Skewness Proxies in Monetary Policy Rules[a]

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
$\sum_{j=1}^{2} a_j$	.69***	.68***	.71***	.70***	.72***	.69***
	(.03)	(.03)	(.04)	(.04)	(.03)	(.03)
β	1.73***	1.55***	2.02***	2.09***	1.74***	1.69***
	(.12)	(.11)	(.16)	(.16)	(.10)	(.12)
γ	.80***	.71***	.80***	.74***	.89***	.81***
	(.06)	(.06)	(.07)	(.08)	(.08)	(.07)
SPD		−.56***				
		(.14)				
sInf			.23**			
			(.10)			
sGDP				−.15		
				(.11)		
DsInf					.40***	
					(.13)	
DsGDP						−.16
						(.12)
LM[b]	.53	.90	.34	.67	.61	.62
Obs.	86	86	68	68	86	86

a. Standard errors are robust to heteroskedasticity. Statistical significance at the ***1, **5, and *10 percent levels.

b. Entries in the "LM" row are *p*-values of Durbin's test for the null hypothesis of no serial correlation in the residuals up to the second order.

point estimates for skewness in the GDP outlook (sGDP and DsGDP) have surprisingly negative signs. However these coefficients are relatively small and insignificant.

Taken together, these results indicate 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. The effects we find suggest that the FOMC acted aggressively to offset concerns about declining growth or rising inflation. We conclude from this econometric analysis that risk management does not just appear in the words of the FOMC—it is reflected in the FOMC's deeds as well.

IV. Conclusion

We have focused on risk surrounding the forecast as a relevant consideration for monetary policy near the ZLB, but other issues are relevant to the liftoff calculus as well. In particular, policymakers may face large reputational 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. Uncertainty over the outlook may be one reason for this persistence. But another reason why policymakers might be reluctant to reverse course is that doing so would damage their reputation, perhaps because the public would lose confidence in the central bank's ability to understand and stabilize the economy. With high uncertainty, this reputational element would lead to more caution. In the case of liftoff, it argues for a longer delay in raising rates to avoid the reputational costs of reverting to the ZLB.

Another reputational concern is the signal the public might infer about the central bank's commitment to its stated policy goals. With regard to liftoff, suppose it occurred with output or inflation still far below target. Large gaps on their own pose no threat to the central bank's credibility if the public is confident that the economy is on a path to achieve its objectives in a reasonable period and that it is willing to accommodate this path. However, if there is uncertainty over the strength of the economy, early liftoff might be construed as a less-than-enthusiastic endorsement of the central bank's ultimate policy objectives. Motivated by the current situation, we have focused in the paper on the case of a central bank that is undershooting its inflation target, but similar issues would arise if risk management considerations dictated an aggressive tightening to guard against inflation and the central bank failed to act accordingly. In a wide class of models, such losses of credibility can have deleterious consequences for achieving the central bank's objectives.

ACKNOWLEDGMENTS We thank numerous seminar participants and Gadi Barlevy, Jeffrey Campbell, Stefania D'Amico, Alan Greenspan, Alejandro Justiniano, John Leahy, Sydney Ludvigson, Leonardo Melosi, Taisuke Nakata, Serena Ng, Valerie Ramey, David Reifschneider, Glenn Rudebusch, Paolo Surico, François Velde, and Johannes Wieland for their help and comments; Theodore Bogusz, David Kelley and Trevor Serrao for superb research assistance; and the volume editors. We also thank Michael McMahon for providing us with machine-readable FOMC minutes and transcripts and Thomas Stark for help with the Philadelphia Fed's real-time data. The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Open Market Committee or the Federal Reserve System.

References

Adam, Klaus and Roberto M. Billi. 2007. "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates." *Journal of Monetary Economics* 54, no. 3: 728–52.

Alcidi, Cinzia, Alessandro Flamini, and Andrea Fracasso. 2011. "Policy Regime Changes, Judgment and Taylor Rules in the Greenspan Era." *Economica* 78: 89–107.

Andrade, Philippe, Eric Ghysels, and Julien Idier. 2013. "Tails of Inflation Forecasts and Tales of Monetary Policy." Kenan-Flagler Research Paper No. 2013-17, University of North Carolina.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2015. "Measuring Economic Policy Uncertainty." Working Paper. http://www.policyuncertainty.com/media/BakerBloomDavis.pdf

Barlevy, Gadi. 2011. "Robustness and Macroeconomic Policy." *Annual Review of Economics* 3: 1–24.

Barsky, Robert, Alejandro Justiniano, and Leonardo Melosi. 2014. "The Natural Rate and Its Usefulness for Monetary Policy Making." *American Economic Review* 104, no. 4: 37–43.

Basu, Susanto, and Brent Bundick. 2014. "Downside Risk at the Zero Lower Bound." Working Paper. http://www.cla.auburn.edu/economics/assets/File/BasuBundickDownsideRisk.pdf

Bekaert, Geert, Marie Hoerova, and Marco Lo Duca. 2013. "Risk, Uncertainty and Monetary Policy." *Journal of Monetary Economics* 60: 771–88.

Bernanke, Ben S. 2012. "The Changing Policy Landscape." In *Monetary Policy since the Onset of the Crisis, Economic Policy Symposium,* pp. 1–22. Kansas City: Federal Reserve Bank of Kansas City.

Bloom, Nicholas. 2009. "The Effect of Uncertainty Shocks." *Econometrica* 77, no. 3: 623–85.

Bomfin, A., and L. Meyer. 2010. "Quantifying the Effects of Fed Asset Purchases on Treasury Yields." Monetary Policy Insights: Fixed Income Focus. Macroeconomics Advisors (online).

Born, Benjamin, and Johannes Pfeifer. 2014. "Policy Risk and the Business Cycle." *Journal of Monetary Economics* 68: 68–85.

Brainard, William. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57, no. 2: 411–25.

Campbell, Jeffrey R., Charles. L. Evans, Jonas D. Fisher, and Alejandro Justiniano. 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity,* Spring, 1–54.

Castelnuovo, Efrem. 2003. "Taylor Rules, Omitted Variables, and Interest Rate Smoothing in the U.S." *Economics Letters* 81, no. 1: 55–59.

Chen, Han, Vasco Curida, and Andrea Ferrero. 2012. "The Macroeconomic Effects of Large Scale Asset Purchase Programmes." *Economic Journal* 122, no. 564: 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, no. 8: 1705–23.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113, no. 1: 1–45.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115, no. 1: 147–80.

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes 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, no. 4: 1371–1406.

Cúrdia, Vasco, Andrea Ferrero, Ging Cee Ng, and Andrea Tambalotti. 2015. "Has U.S. Monetary Policy Tracked the Efficient Interest Rate." *Journal of Monetary Economics* 70 (March): 72–83.

D'Amico, Stefania, and Thomas King. 2013. "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply." *Journal of Financial Economics* 108, no. 2: 425–448.

———. 2015. "Policy Expectations, Term Premia, and Macroeconomic Performance." Unpublished manuscript. Chicago: Federal Reserve Bank of Chicago.

D'Amico, Stefania, and Athanasios Orphanides. 2014. "Inflation Uncertainty and Disagreement in Bond Risk Premia." Working Paper 2014-24. Chicago: Federal Reserve Bank of Chicago.

Dolado, Juan J., P. Ramón María-Dolores, and Manuel Naveira. 2005. "Are Monetary Policy Reaction Functions Asymmetric? The Role of Nonlinearity in the Phillips Curve." *European Economic Review* 49, no. 2: 485–503.

Dolado, Juan J., P. Ramón María-Dolores, and Francisco J. Ruge-Murcia. 2004. "Nonlinear Monetary Policy Rules: Some New Evidence for the U.S." *Studies in Nonlinear Dynamics and Econometrics* 8, no. 3.

Egertsson, Gauti B., and Michael Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 2003, no. 1: 139–211.

Eichenbaum, Martin, and Jonas D. Fisher. 2007. "Estimating the Frequency of Price Re-Optimization in Calvo-Style Models." *Journal of Monetary Economics* 54, no. 7: 2032–47.

Engen, Eric M., Thomas T. Laubach, and David Reifschneider. 2015. "The Macroeconomic Effects of the Federal Reserve's Unconventional Monetary Policy." Finance and Economics Discussion no. 2012-5. Washington: Board of Governors of the Federal Reserve System.

English, William B., J. David Lopez-Salido, and Robert Tetlow. 2013. "The Federal Reserve's Framework for Monetary Policy—Recent Changes and New Questions." Finance and Economics Discussion no. 2013-76. Washington: Board of Governors of the Federal Reserve System.

Evans, Charles L. 2014. "Patience Is a Virtue When Normalizing Monetary Policy." Speech presented at the Conference on Labor Market Slack, Peterson Institute for International Economics, Washington, September 24.

Fernandez-Villaverde, Jesus, Pablo Guerron-Quintana, Keith Kuester, and Juan Rubio-Ramirez. 2012. "Fiscal Volatility Shocks and Economic Activity." Working Paper no. 11-32/R. Philadelphia: Federal Reserve Bank of Philadelphia.

Friedman, Benjamin M. 1975. *Economic Stabilization Policy: Methods in Optimization.* Amsterdam and New York: North-Holland Publishing Company and American Elsevier Publishing Company.

Fuhrer, Jeffrey C. 2000. "Habit Formation in Consumption and Its Implications for Monetary-Policy Models." *American Economic Review* 90, no. 3: 367–90.

Gagnon, J., M. Raskin, J. Remache, and B. Sack. 2010. "Optimal Fiscal and Monetary Policy with Occasionally Binding Zero Bound Constraints." Staff Report no. 441. New York: Federal Reserve Bank of New York.

Galí, Jordi. 2008. *An Introduction to the New Keynesian Framework.* Princeton University Press.

Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44, no. 2: 192–222.

Gerlach-Kristen, Petra. 2004. "Interest Rate Smoothing: Monetary Policy Inertia or Unobserved Variables." *Contributions in Macroeconomics* 4, no. 1: 1534–6005.

Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102, no. 4: 1692–1720.

Gnabo, Jean-Yves, and Diego N. Moccero. 2015. "Risk Management, Nonlinearity and Aggressiveness in Monetary Policy: The Case of the U.S." *Journal of Banking & Finance* 55: 281–94.

Goodfriend, Marvin. 1991. "Interest Rates and the Conduct of Monetary Policy." *Carnegie-Rochester Conference Series on Public Policy* no. 34: 7–30.

Greenspan, Alan. 2004. "Risk and Uncertainty in Monetary Policy." *American Economic Review* 94, no. 2: 33–40.

Hamilton, James D., Ethan S. Harris, Jan Hatzius, and Kenneth D. West. 2015. "The Equilibrium Real Funds Rate: Past, Present and Future." Unpublished manuscript, University of California—San Diego.

Hamilton, James, and Jing Wu. 2010. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." Unpublished manuscript, University of California—San Diego.

Hansen, Lars, and Thomas Sargent. 2008. "Robustness." Princeton University Press.

Jurado, Kyle, Sydney Ludvigson, and Serena Ng. 2015. "Measuring Uncertainty." *American Economic Review* 105, no. 3: 1177–1276.

Kiesel, Konstantin, and Maik H. Wolters. 2014. "Estimating Monetary Policy Rules When the Zero Lower Bound on Nominal Interest Rates Is Approached." Kiel Working Paper no. 1898, Kiel Institute for the World Economy.

Kiley, Michael T. 2012. "The Aggregate Demand Effects of Short- and Long-Term Interest Rates." Finance and Economics Discussion no. 2012-54. Washington: Board of Governors of the Federal Reserve System.

Kilian, Lutz, and Simone Manganelli. 2008. “The Central Banker as a Risk Manager: Estimating the Federal Reserve’s Preferences under Greenspan.” *Journal of Money, Credit and Banking* 40, no. 6: 1103–29.

Krishnamurthy, Arvind, and Annette Vissing-Jørgensen. 2013. “The Ins and Outs of LSAPS.” *Proceedings of the Economic Policy Symposium—Jackson Hole.* Federal Reserve Bank of Kansas City.

Krugman, Paul R. 1998. “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap.” *Brookings Papers on Economic Activity,* no. 2: 137–87.

Laubach, Thomas, and John C. Williams. 2003. “Measuring the Natural Rate of Interest.” *Review of Economics and Statistics* 85. no. 4: 1063–70.

Laxton, Douglas, David Rose, and Demosthenes Tambakis. 1999. “The U.S. Phillips Curve: The Case for Asymmetry.” *Journal of Economic Dynamics & Control* 23, nos. 9, 10: 1459–85.

Mas-Colell, Andreau, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory.* Oxford University Press.

Mumtaz, Haroon, and Paolo Surico. 2015. “The Transmission Mechanism in Good and Bad Times.” *International Economic Review,* forthcoming.

Nakata, Taisuke. 2013a. “Optimal Fiscal and Monetary Policy with Occasionally Binding Zero Bound Constraints.” Finance and Economics Discussion no. 2013-40. Washington: Board of Governors of the Federal Reserve System.

———. 2013b. “Uncertainty at the Zero Lower Bound.” Finance and Economics Discussion no. 2013-09. Washington: Board of Governors of the Federal Reserve System.

Nakata, Taisuke, and Sebastian Schmidt. 2014. “Conservatism and Liquidity Traps.” Finance and Economics Discussion no. 2014-105. Washington: Board of Governors of the Federal Reserve System.

Nakov, Anton A. 2008. “Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate.” *International Journal of Central Banking* 4, no. 2: 73–127.

Orphanides, Athanasios, and John C. Williams. 2002. “Robust Monetary Policy Rules with Unknown Natural Rates.” *Brookings Papers on Economic Activity,* Fall: 63–145.

Reifschneider, David, and John C. Williams. 2000. “Three Lessons for Monetary Policy in a Low-Inflation Era.” *Journal of Money, Credit, and Banking* 32, no. 4: 936–66.

Romer, Christina D., and David H. Romer. 1989. “Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz.” In *NBER Macroeconomics Annual 2005,* vol. 4, ed. by O. Blanchard and S. Fischer.

Rudebusch, Glenn D. 2002. “Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia.” *Journal of Monetary Economics* 49, no. 6: 1161–87.

Rudebusch, Glenn, and Lars E. Svensson. 1999. “Policy Rules for Inflation Targeting.” In *Monetary Policy Rules,* ed. by J. B. Taylor. University of Chicago Press.

Sack, Brian. 2000. “Does the Fed Act Gradually? A VAR Analysis.” *Journal of Monetary Economics* 46, no. 1: 229–56.

Smets, Frank, and Rafael Wouters. 2007. “Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach.” *American Economic Review* 97, no. 3: 586–606.

Surico, Paolo. 2007. “The Fed’s Monetary Policy Rule and U.S. Inflation: The Case of Asymmetric Preferences.” *Journal of Economic Dynamics & Control* 31, no. 1: 305–24.

Svensson, Lars, and Michael Woodford. 2002. “Optimal Policy with Partial Information in a Forward-Looking Model: Certainty-Equivalence Redux.” Working Paper, Columbia University. http://www.columbia.edu/~mw2230/SWCE206.pdf

———. 2003. “Indicator Variables for Optimal Policy.” *Journal of Monetary Economics* 50, no. 3: 691–720.

Taylor, John B. 1993. “Discretion versus Policy Rules in Practice.” Carnegie-Rochester Conference Series on Public Policy no. 39: 195–214.

Tenreyro, Silvana, and Gregory Thwaites. 2015. “Pushing on a String: US Monetary Policy Is Less Powerful in Recessions.” Working Paper, London School of Economics and Political Science. http://personal.lse.ac.uk/tenreyro/TandT.pdf

Werning, Iván. 2012. “Managing a Liquidity Trap: Monetary and Fiscal Policy.” Unpublished manuscript, Massachusetts of Technology.

Woodford, Michael. 2003. “Interest and Prices.” Princeton University Press.

———. 2012. “The Changing Policy Landscape.” In *Methods of Policy Accommodation at the Interest-Rate Lower Bound, Economic Symposium* (pp. 185–288). Kansas City: Federal Reserve Bank of Kansas City.