PRELIMINARY AND INCOMPLETE FIRST DRAFT

Risk Management for Monetary Policy at the Zero Lower Bound*

Charles Evans   Jonas Fisher   François Gourio   Spencer Krane

January 23, 2015

Abstract

As labor markets improve and projections have inflation heading back toward target, the Fed has begun to contemplate lifting the federal funds rate from its zero lower bound (ZLB). Under what conditions should the Fed start raising rates? We lay out an argument that calls for caution. It is founded on a *risk management* principle that says policy should be formulated taking into account the dispersion of outcomes around the mean forecast. On the one hand, raising rates early increases the likelihood of adverse shocks driving a fragile economy back to the ZLB. On the other hand, delaying lift-off when the economy turns out to be resilient could lead to an unwelcome bout of inflation. Since the tools available to counter the first scenario are hard to implement and may be less effective than the traditional tool of raising rates to counter the second scenario, the costs of premature lift-off exceed those of delay. This article shows in a canonical framework that uncertainty about being constrained by the ZLB in the future implies an optimal policy of delayed lift-off. We present evidence that such a risk management policy is consistent with past Fed actions and that unconventional tools will be hard to implement if the economy were to be constrained by the ZLB after a hasty exit.

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

*All the authors are affiliated with the Federal Reserve Bank of Chicago. We thank Gadi Barlevy, Jeffrey Campbell, Stefania D'Amico, Alejandro Justiniano, Leonardo Melosi and François Velde for helpful comments, and Theodore Bogusz, David Kelley and Trevor Serrao for superb research assistance. 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.

1 Introduction

Current forecasts by most Federal Open Market Committee (FOMC) participants look for the unemployment rate to return to its long-run neutral level by 2016 and for inflation to gradually rise back to its 2 percent target. This forecast could go wrong in two ways. One is that the FOMC may be overestimating the underlying strength in the real economy. Guarding against this risk calls for a more patient removal of accommodation. The second is that the FOMC may be wrong about the inflation outlook, and the economy could be poised for a much stronger rise in inflation than is currently projected. This risk calls for more aggressive rate hikes. How should policy should manage these divergent risks?

The biggest risk we face today is prematurely engineering restrictive monetary conditions. If the FOMC misjudges the impediments to growth and reduces monetary accommodation too soon, it could find itself in the very uncomfortable position of falling back into the ZLB environment. The implications of the ZLB for growth and employment are severe. It is true the FOMC has access to unconventional policy tools while at the ZLB, but their is no guarantee they will be as successful as they have been in the past if the economy were to return after a brief exit. It seems likely that the credibility underlying the prior use of these policies will be diminished by an unduly hasty exit from the ZLB. Furthermore, building a consensus among the Committee to use them may be harder to come by. And there is also a non-trivial possibility that their potency has been substantially diminished. These considerations suggest that unconventional tools will be weaker substitutes for the conventional interest rate tool, placing the attainment of the FOMC's mandated goals at greater risk.

In contrast, it is reasonable to imagine that the costs of inflation running moderately above target for a while are much smaller than the costs of falling back into the ZLB. This is not the least because it is likely that inflation could be brought back into check with modest increases in interest rates. These measured rate increases likely would be manageable for the real economy, particularly if industry and labor markets had already overcome the headwinds that have kept productive resources from being efficiently and fully employed. In addition, inflation in the U.S. has averaged well under that 2 percent mark for the past six and a half years. With a symmetric inflation target, one could imagine moderately-above-target inflation for a limited period of time as simply the flip side of the recent inflation experience — and hardly an event that would impose great costs on the economy.

To summarize, raising rates early increases the likelihood of adverse shocks driving a fragile economy back to the ZLB. But, delaying lift-off when the economy turns out to be resilient could lead to an unwelcome bout of inflation. Since the tools available to counter the first scenario are hard to implement and may be less effective than the traditional tool of raising rates to counter the second scenario, the costs of premature lift-off exceed those of delay. It therefore seems prudent to delay lift-off from the ZLB as long as possible to be sure that the economy has regained its resilience and a sustained period of strong growth is underway.

In this paper we establish within a canonical framework that uncertainty about monetary policy being constrained by the ZLB in the future implies an optimal policy of delayed lift-off. This result rationalizes the risk management policy just described. We define risk management in general as the principle that policy should be formulated taking into account the dispersion of outcomes around the mean forecast. In a wide class of models, and in particular the New Keynesian model, optimal policy involves adjusting the interest rate in response to mean forecasts alone, and information on the dispersion of the forecast, such its variance or skewness, is irrelevant. This is the certainty equivalence principle. We review the extensive literature that demonstrates how departures from the canonical framework, such as asymmetric central bank preferences or uncertainty about the effects of interest rate changes on the economy, yield roles for risk management *away* from the ZLB, *i.e.* departures from certainty equivalence. Our main theoretical contribution is to demonstrate that within the canonical framework the ZLB implies a new role for risk management that is otherwise not present.

We view the ZLB as a key source of asymmetry in the standard New Keynesian model; this asymmetry naturally generates a risk management concern. However, the intuition for the result turns out to be subtle. We first show that uncertainty about whether the ZLB binds *today* does not, in itself, generate a motive for looser policy - the certainty equivalence principle holds despite the constraint. Uncertainty about whether the ZLB binds in the *future* does matter, however. We present in detail a normative analysis of two distinct economic mechanisms that make optimal policy under discretion respond to risk due to the ZLB. We stress that these mechanisms operate in very standard macroeconomic models, so that no leap of faith is necessary to embrace these results, at least at a qualitative level.

The first channel - which we call the *expectations channel* – arises because the possibility of a binding ZLB tomorrow leads to lower expected inflation and output today, and hence requires some policy easing. The second channel - which we call the *buffer stock* channel – arises because it can be useful to build up output or inflation today in order to reduce the likelihood of hitting the ZLB tomorrow, or at least to reduce the severity of the potential ZLB episode tomorrow. We show that optimal policy when one or both of these channels are operative is qualitatively very similar to the policy described above. In particular optimal policy dictates that lift-off from a zero interest rate at a time when a return to the ZLB remains a possibility should be tempered.

While we establish a solid theoretical basis for a risk management approach at the ZLB it is natural to ask whether proposing risk management in the current policy environment would be a departure from how the FOMC has acted in the past. Is our proposal something out of the ordinary other than the fact that we are at the ZLB? We explore this question in two ways.

First, using our review of optimal policy away from the ZLB as a guide, we analyze the FOMC minutes and other monetary policy communications and find evidence that risk management has been a long-standing operating characteristic of the FOMC, at least in words if not in deeds. We find that there are numerous examples when uncertainty and insurance have been used to explain monetary policy settings. This analysis demonstrates that calling for a risk management approach in the current policy environment is not out of the ordinary and in fact is a well-established approach to monetary policy. Confirmation of this view is found in Greenspan (2004) who states "...the conduct of monetary policy in the United States has come to involve, at its core, crucial elements of risk management."

Second, we explore whether the words of the FOMC are reflected in policy actions. For this analysis we estimate a conventionally specified policy reaction function and investigate whether the coefficients on a variety of measures of risk are significantly different from zero when added to the reaction function. The measures of risk we look at include ones based on financial market data, survey measures of forecasts, and several measures derived from our analysis of the FOMC minutes. While the findings are not overwhelming there is certainly evidence that risk in the economic outlook has had a material impact on the interest rate choices of the FOMC. This work complements a well-established literature that estimates non-linear policy reaction functions and other work that has investigated empirically the role of risk in monetary policy.

We also address the key component of our policy proposal that unconventional policies at the ZLB are not perfect substitutes for conventional policies away from it. Our theoretical analysis assumes that the only instrument available to the policy-maker is the short term nominal interest rate. Yet the FOMC has used unconventional tools such as forward guidance and large scale asset purchases (LSAPs). If these tools are virtually equivalent to the interest rate tool then from a theoretical perspective the ZLB implies no special role for risk management.

We argue that unconventional tools are imperfect substitutes primarily by studying speeches by Fed officials, dealer surveys, and special questions in the Blue Chip survey. This study establishes that there are widespread doubts about the efficacy of unconventional policies. Furthermore resistance within the FOMC could impose barriers to their use in the future. We also review evidence showing that LSAPs have impacted the economy in large part through signalling that short term rates will remain low for a considerable period; they reinforce forward guidance. Therefore any benefit must have derived from the credibility of the FOMC to follow through. A hasty exit from the ZLB would put this credibility at risk thereby limiting the effectiveness of unconventional policies should the economy be constrained by the ZLB again.[1]

2 Theoretical Rationales for Risk Management at the ZLB and Beyond

The canonical framework of monetary policy analysis assumes that the central bank sets the nominal interest rate to minimize a quadratic loss function of inflation and the output gap, and that the economy is described by a set of linear equations. This framework allows to calculate the optimal interest rate, as a function of given "economic fundamentals." 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.[2] In most applications, uncertainty is incorporated as additive shocks to these linear equations capturing factors outside the model that lead to variation in economic activity or inflation. Under these conditions, a general result is the so-called "certainty equivalent principle": the extent of uncertainty does not affect the optimal interest rate.[3] Moreover, the optimal interest rate is a linear function of the fundamentals, which implies that the policymaker should react (at the margin) in the same way to fluctuations in fundamentals - be they small or large, negative or positive.

The certainty equivalence feature inherent to the canonical framework is analytically quite convenient. It allows us to cut through technical difficulties to obtain an intuition about the role of policy. However, a limitation of this approach is that by construction, it denies that a policymaker might choose to adjust his or her policy in the face of an increase in uncertainty that leaves forecasts unchanged (*i.e.* mean-preserving spreads in the distribution of shock terms). However, anecdotal and empirical evidence discussed below in Section 3

[1] The work described in this paragraph will appear in the next draft.

[2] For instance, the Board staff routinely conducts optimal policy exercises in the FRB-US model, see English, López-Salido, and Tetlow (2013) or http://www.federalreserve.gov/econresdata/notes/feds-notes/2014/optimal-control-monetary-policy-in-frbus-20141121.html.

[3] A general statement of the certainty equivalence principle can be found in Chow (1973). A more recent restatement in models with forward-looking variables is in Svensson and Woodford (2002, 2003). Note that this result does not require that shocks are Gaussian or identically distributed over time. (One exception is when some states are unobserved by agents and/or the policy maker and the Kalman filter is used for signal extraction. To preserve the overall linearity of equations then requires gaussian shocks.)

suggests that in practice, policymakers are sensitive to uncertainty and respond accordingly, a custom sometimes referred to as a risk management approach to policy. To understand why central bankers behave this way requires some departure from the canonical framework, e.g. non-quadratic loss functions, non-linearities in the decisions economic agents make, or non-additive shocks (e.g. shocks to the coefficients that govern the behavior of economic variables). Previous work has explored these; we review this literature below in detail.

The main contribution of this section, however, is to consider a different departure from the canonical model associated with the possibility of a binding zero lower bound (ZLB). That is, we will argue that once we acknowledge the possibility that within the canonical model a policymaker will ever be constrained by an inability to set negative nominal rates (or a limit on how negative they might be), certainty equivalence breaks down and optimal policy should take uncertainty into account. In particular, we show that an increase in uncertainty that raises the odds of the ZLB binding at some point in the future should lead a central bank to reduce its current policy rate. An implication of this is that liftoff from a zero interest rate at a time when a return to the ZLB remains a possibility should be tempered.

This section is organized as follows. We first present our baseline economic model, and use it to illustrate why uncertainty is irrelevant if the uncertainty is about whether the ZLB binds *today*. Next, we analyze the two economic channels that justify shading policy towards lower rates when the likelihood of a *future* ZLB becomes significant. Finally, we review other mechanisms that can be relevant for risk management even when the policymakers believe there is no risk that the ZLB will bind in the future. These mechanisms have been studied in the extant literature, and they may be important in accounting for the patterns that we document in our narrative and empirical sections below.

2.1 Assumptions and certainty equivalence results

We start by describing assumptions that we make for most of the theory analysis, then we discuss standard optimal policy results and illustrate that uncertainty about the ZLB today does not, in itself break the certainty equivalence result in this model.

2.1.1 Assumptions

For most of this paper, we use the standard forward-looking New Keynesian model. Given that there are many excellent expositions of this model, e.g. Woodford (2003) or Gali (2008), we just state our notation without much explanation. The model consists of two main equations, the so-called IS (or consumer Euler equation) curve and Phillips curves. The Phillips curve reads

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

where π_t is inflation, x_t is the output gap, and u_t is a "cost-push" shock with mean zero, β is the discount factor, and κ measures the sensitivity of inflation to the output gap and plays a key role in optimal policy since it affects the trade-off between inflation and output. The IS curve reads,

$$y_t = E_t y_{t+1} - \frac{1}{\sigma}\left(i_t - \rho - E_t \pi_{t+1}\right) + g_t,$$

where y_t is (log) output, i_t the nominal interest rate (controlled by the central bank), g_t a demand shock with mean zero, and σ is the elasticity of intertemporal substitution. To link the output gap to output, we define (log) potential as z_t, which is assumed to follow an exogenous process. The output gap is then given by $x_t = y_t - z_t$, so that we can rewrite the IS curve as:

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

where ρ_t^n is the so-called "natural rate of interest",

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

In what follows, we will often make simplifying assumptions on the structure of shocks, and in particular assume that there are no shocks after a certain time, *i.e.* $\rho_t^n = \rho > 0$ and $u_t = 0$ for $t \geq T$; this allows an easy characterization of the "long-term" values of output and inflation, which then allows to work backwards in time the value of output and inflation at all previous dates $t < T$. Last, note that this model is written in deviations from a steady-state, so negative values mean "below target".

Turning to the assumptions about how interests rate are set, we follow the canonical framework in assuming a quadratic loss function, with weight λ respectively on output gap deviations:

$$L = \frac{1}{2} E \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda x_t^2 \right).$$

We further assume that there is a hard bound on the nominal interest rate, $i_t \geq 0$.[4]

Most of our analysis will solve for optimal policy under discretion; that is, each period the central bank sets the nominal interest rate, given the situation today, and private agents anticipate that the central bank will re-optimize tomorrow. We believe all our results hold qualitatively if policy was instead set under full commitment, as we discuss later.[5]

2.1.2 A simple illustration of optimal policy

To warm up, we start by presenting some standard results regarding optimal monetary policy in the simple model described above.[6] Suppose the central bank sets the interest rate at time 0 after observing the natural rate ρ_0^n and the cost-push shock u_0. And suppose for now that it is known both to private agents and to the central bank that there will not be shocks from time $t = 1$ on, a blunt way to describe that the economy will "return to normal" tomorrow; *i.e.* $u_t = 0$ and $\rho_t^n = \rho$ for $t \geq 1$. This implies that, starting at time 1, the central bank can reach the best possible outcome, which is $x_t = \pi_t = 0$, by setting the interest rate $i_t = \rho$.

[4] While some central banks, such as the European Central Bank (ECB) and the Swiss National Bank (SNB) have been able to introduce slightly negative nominal rates, there is clearly a limit to how negative the nominal rate can be, at least under current institutional arrangements, before savers turn to cash. Hence, while the true bound might not be exactly zero, it is likely to be some small negative number. It is true that other tools - such as forward guidance or LSAPs - can be used if the ZLB significantly constrains policy. We argue in Section 5 that if the economy were to return to the ZLB after a hasty exit, these tools may not be very effective. So we ignore them here.

[5] It is well known from the contributions of Krugman (1998), Egertsson and Woodford (2003), Woodford (2012) and Werning (2012) that commitment can reduce markedly the severity of the ZLB problem by creating higher expectations of inflation and the output gap.

[6] Clarida, Gali, and Gertler (1999) present these results (and several more) in a more general framework.

The problem of choosing the interest rate at time 0 then becomes a simple static problem:

$$
\begin{aligned}
& \min_{i_0} \frac{1}{2}\left(\pi_0^2+\lambda x_0^2\right) \\
s.t. \quad & : \\
\pi_0 &= \kappa x_0+u_0, \\
x_0 &= -\frac{1}{\sigma}\left(i_0-\rho_0^n\right), \\
i_0 &\geq 0.
\end{aligned}
$$

where we have used the fact that expected time 1 inflation and output are zero, $E\pi_1 = Ex_1 = 0$. Next, note that we can rewrite the ZLB constraint as

$$x_0 \leq \frac{\rho_0^n}{\sigma},$$

and the choice of an interest rate is equivalent to a choice of an output gap. The solution to this problem is simple. First, calculate the solution assuming that the ZLB constraint does not bind:

$$\min_{x_0} \frac{1}{2}\left(\left(\kappa x_0+u_0\right)^2+\lambda x_0^2\right),$$

leading to

$$x_0=-\frac{\kappa}{\lambda+\kappa^2} u_0,$$

and hence

$$
\begin{aligned}
\pi_0 &= \frac{\lambda}{\lambda+\kappa^2} u_0, \\
i_0 &= \rho_0^n+\frac{\sigma\kappa}{\lambda+\kappa^2} u_0.
\end{aligned}
$$

This solution features two standard results that hold true if the ZLB constraint does not bind: (1) shocks to the natural rate of interest ρ_0^n do not affect the output gap or inflation since they can be perfectly offset by an appropriate choice of the nominal rate; (2) cost-push shocks, however, generate a trade-off between inflation and output; the central bank responds by offsetting partially the effect of such shocks on inflation, which generates an opposite movement in output; the optimal response depends on the slope of the Phillips curve κ and the weight the central bank puts on output stabilization.

This solution, however, may not be feasible if the interest rate required to implement it is negative, *i.e.* if

$$\rho_0^n + \frac{\sigma\kappa}{\lambda + \kappa^2} u_0 \leq 0,$$

which occurs if the realizations of the demand shock and/or the cost-push shocks are too low. In this case, the central bank cannot lower the interest rate enough to engineer the output boom it needs to offset either the demand or cost-push shocks. The optimal solution then involves simply setting the interest rate $i_0 = 0$, which yields

$$\begin{aligned} x_0 &= \frac{\rho_0^n}{\sigma}, \\ \pi_0 &= \kappa\frac{\rho_0^n}{\sigma} + u_0. \end{aligned}$$

When the ZLB binds, the economy has lower output and lower inflation that would otherwise be the case. In that case, (1) demand shocks affect the output gap since monetary policy cannot respond, and (2) cost-push shock fully go through to inflation.

2.1.3 Why uncertainty about the ZLB today does not matter

To illustrate how uncertainty affects this standard monetary policy analysis, assume that the central bank has to set the interest rate before seeing the shocks. This timing protocol captures that it is difficult for the central bank to observe the economy in real-time, for a variety of reasons.[7] The problem solved by the central bank is to then minimize the *expected* loss,

$$\min_{i_0} \frac{1}{2} E\left(\pi_0^2 + \lambda x_0^2\right),$$

[7]For now we assume still that it is known for sure that the economy will exit the ZLB tomorrow, with $\rho_t^n = \rho > 0$ and $u_t = 0$ for all $t \geq 1$.

$$
\begin{aligned}
s.t. \quad &: \\
\pi_0 &= \kappa x_0 + u_0, \\
x_0 &= -\frac{1}{\sigma}\left(i_0 - \rho_0^n\right), \\
i_0 &\geq 0.
\end{aligned}
$$

With some simple algebra, we can see that the solution to this problem is the interest rate that solves[8]

$$
\begin{aligned}
&\min_{i_0} \frac{\kappa^2 + \lambda}{2\sigma^2} E(i_0 - \rho_0^n)^2, \\
s.t. \quad : \quad &i_0 \geq 0.
\end{aligned}
$$

This is a standard root-mean squared error problem, with solution $i_0 = E(\rho_0^n)$, if $E(\rho_0^n) > 0$; and $i_0 = 0$ if $E(\rho_0^n) \leq 0$. This result illustrates the certainty equivalence property: when faced with uncertainty about the "supply" (u_0) and "demand" (ρ_0^n) shocks, the central bank will set its interest rate according to its mean forecasts of the shocks.[9] In particular, the amount of uncertainty about ρ_0^n or u_0 does not affect the decision of the policymaker. (It does, however, impact the expected loss – the central bank dislikes uncertainty, since it makes it harder to achieve its goals on average.) This result shows that uncertainty regarding whether the ZLB binds *today* does not, in itself, lead to a violation of the certainty equivalent principle. However, as we now will discuss, the possibility of a binding ZLB *tomorrow* will turn out to affect optimal policy *today*.

[8] The loss function includes some other terms, but these are independent of the value of i_0 given our assumption that $E(u_0) = 0$. Specifically,

$$
\frac{1}{2}E(\pi_0^2 + \lambda x_0^2) = \frac{\kappa^2 + \lambda}{2\sigma^2}E(i_0 - \rho_0^n)^2 + \frac{1}{2}E(u_0^2) + \frac{1}{\sigma}E(u_0\rho_0^n).
$$

[9] We assumed $E(u_0) = 0$. If $E(u_0) \neq 0$, then the optimal nominal interest rate is given by the formula:

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

if this quantity is positive, and zero otherwise.

2.2 The expectations channel

This section describes the first channel through which the possibility of a future binding ZLB affects optimal policy today. We first describe a parsimonious macroeconomic scenario that allows us to discuss the role of ZLB uncertainty while maintaining tractability. We then solve for optimal policy given this scenario. Finally we provide some simple examples and discuss the implications and limitations of our analysis.

2.2.1 A ZLB uncertainty scenario

We assume that the central bank observes the current value of the natural rate of interest, ρ_0^n, as well as the cost-push shock u_0; moreover, there is no uncertainty after time 2, specifically $\rho_t^n = \rho > 0$ and $u_t = 0$ for all $t \geq 2$. However, there is uncertainty at time 1 regarding either the level of the natural rate of interest ρ_1^n or the cost-push shock u_1. (In the interest of simplicity, we analyze the uncertainty about the two shocks separately; analyzing them jointly complicates the notation without providing any additional insight.) The variables ρ_1^n and u_1 are assumed to be distributed according to probability density functions $f_\rho(.)$ and $f_u(.)$ respectively. This very simple specific stochastic structure keeps the optimal policy calculation tractable while preserving the main insights, and we think it captures some key elements of uncertainty today. First, the economy will eventually exit the ZLB regime. Second, there is uncertainty as to when this will happen: the natural rate might be low enough at time 1 that the ZLB binds – a "delayed recovery" scenario; or the natural rate of interest may be high at time 1 as the economy recovers, which could potentially lead to some inflation.[10]

A word is in order regarding the interpretation of these shocks. The natural rate of interest would likely capture factors affecting demand such as fiscal policy, foreign economies growth, and financial factors such as deleveraging in the United States. The cost-push

[10] We do not for now take a stance on whether the ZLB is still binding today - *i.e.*, the values of the shocks ρ_0^n and u_0 might or not be low. The ZLB will endogenously bind at time 0 for some parameter values, as we will see.

shock could capture factors affecting inflation such as an independent decline in inflation expectations, as well as other short-run factors such as the dollar appreciation or the decline in oil prices.

Last, we do not need to describe the conditions preceding time 0, since they are not relevant for optimal policy today (technically, because our model has no endogenous state variables), but one possibility is that the natural rate ρ_t^n was negative for $t < 0$ and that the policy rate was set at zero, $i_t = 0$ for $t < 0$, so that the economy is currently close to exiting a period where the ZLB constraint was binding. The model is thus pertinent to discuss the timing of liftoff.

2.2.2 Optimal policy with natural rate uncertainty only

Before delving into the calculations, it may be useful to provide the overall intuition. Private agents look to the future and envision two possibilities - either the ZLB will bind tomorrow, or it will not. If it doesn't, then output and inflation will be zero. But if the ZLB binds, output and inflation will be negative (*i.e.* below target). As a result, overall expected inflation and expected output are lower next period due to the possibility of a binding ZLB. This implies that agents today have lower desired spending and set lower prices (lower inflation today), which requires accommodation today.

In the interest of clarity, we start by solving the model assuming that there are no cost-push shocks at time 1; hence, the only uncertainty pertains to the natural rate of interest. To find the optimal policy under discretion, one can solve the model backwards. First, for $t \geq 2$, it is possible to perfectly stabilize the economy by setting the nominal interest rate equal to the (now positive) natural rate, $i_t = \rho_t^n = \rho$. This leads to $\pi_t = x_t = 0$ for $t \geq 2$.[11] Going backward in time, the optimal policy at time 1 will depend on the realized value of

[11] We note that this simple interest rate rule is consistent with 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, for instance by introducing in the interest rate rule a term that responds to inflation, $i_t = \rho_t^n + \phi\pi_t$. In equilibrium, $\pi_t = 0$, so this term is inoperative, but the threat of high interest rates if inflation is nonzero is important to "kill" other equilibria. From now on, we will not mention this issue. See for instance (Gali 2008, pp. 76–77) for a discussion.

the natural rate ρ_1^n . As in the example of section 2.1, two cases arise. 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 $x_1 = \frac{\rho_1^n}{\sigma} < 0$, and $\pi_1 = \kappa \frac{\rho_1^n}{\sigma} < 0$.

The probability that the ZLB binds is $P = \int_{-\infty}^{0} f_\rho(\rho) d\rho$,and the expected loss is $L = \frac{\lambda+\kappa^2}{2\sigma^2} \int_{-\infty}^{0} \rho^2 f_\rho(\rho) d\rho$. The expected output gap before time 1 is hence $Ex_1 = \frac{1}{\sigma} \int_{-\infty}^{0} \rho f(\rho) d\rho < 0$ and expected inflation is $E\pi_1 = \kappa Ex_1 < 0$.

Because agents are forward-looking, this low expected output gap and inflation at time 1 feed backward to time 0. A low output tomorrow depresses output today by intertemporal substitution. And a low inflation tomorrow depresses inflation today as price setting is forward looking, and depresses output by raising the real interest rate. The optimal policy at time 0 must take into account these effects. This implies that optimal policy will be looser than if there was no chance that the ZLB binds tomorrow. Formally, optimal policy at time 0 minimizes the current loss (since the future expected loss L is independent of policy today):

$$
\begin{aligned}
& \min_{i_0} \frac{1}{2} \left(\pi_0^2 + \lambda x_0^2 \right), \\
s.t. \quad & : \\
x_0 \quad & = \quad Ex_1 - \frac{1}{\sigma} \left(i_0 - \rho_0^n - E\pi_1 \right), \\
\pi_0 \quad & = \quad \beta E\pi_1 + \kappa x_0 + u_0. \\
i_0 \quad & \geq \quad 0.
\end{aligned}
$$

We can rewrite this problem as:

$$
\begin{aligned}
& \min_{x_0} \frac{1}{2} \left((\kappa x_0 + \beta E\pi_1)^2 + \lambda x_0^2 \right), \\
s.t. \quad & : \\
x_0 \quad & \leq \quad Ex_1 + \frac{1}{\sigma} \left(\rho_0^n + E\pi_1 \right).
\end{aligned}
$$

Two cases arise again at this state, depending on whether the ZLB binds at time 0 or not.

Define the threshold value

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

If $\rho_0^n > \rho_0^*$, then the optimal policy is to follow the standard monetary policy response to the inflation disturbance $\beta E\pi_1 + u_0$, as in section 2.1.:

$$\begin{aligned} x_0 &= -\frac{\kappa}{\lambda+\kappa^2} \left(\beta E\pi_1 + u_0 \right), \\ \pi_0 &= \frac{\lambda}{\lambda+\kappa^2} \left(\beta E\pi_1 + u_0 \right). \end{aligned}$$

A full solution is obtained by substituting in $E\pi_1 = \frac{\kappa}{\sigma} \int_{-\infty}^{0} \rho f(\rho) d\rho$. The possibility of ZLB means lowers $E\pi_1$, hence higher output and lower inflation today, that is a looser policy. The corresponding interest rate is

$$\begin{aligned} i_0 &= \rho_0^n + E\pi_1 + \sigma(Ex_1 - x_0), \\ &= \rho_0^n + \sigma \frac{\kappa}{\lambda+\kappa^2} u_0 + \left(\kappa + \sigma + \sigma\beta \frac{\kappa^2}{\lambda+\kappa^2} \right) Ex_1. \end{aligned}$$

Because $Ex_1 < 0$, this formula implies that the optimal interest rate is lower than what would be the case if there was no chance of a ZLB situation tomorrow. The interest rate is lowered to offset the deflationary and recessionary effects of a *potential* ZLB tomorrow. If

$\rho_0^n < \rho_0^*$, then the ZLB binds today and optimal policy amounts to $i_0 = 0$, leading to

$$\begin{aligned} x_0 &= \left(1 + \frac{\kappa}{\sigma} \right) Ex_1 + \frac{\rho_0^n}{\sigma}, \\ \pi_0 &= Ex_1 \left(\beta\kappa + 1 + \frac{\kappa}{\sigma} \right) + \kappa \frac{\rho_0^n}{\sigma} + u_0. \end{aligned}$$

Note that this implies that for some parameters, the ZLB will bind today, at time 0, even though it would not bind if agents were certain that the economy would perform well tomorrow. Specifically, if people were certain that the ZLB would not bind tomorrow, then

$Ex_1 = E\pi_1 = 0$ and the ZLB would only bind at time 0 if $\rho_0^n < \rho^{**} = -\sigma\frac{\kappa}{\lambda+\kappa^2}u_0$, as demonstrated in section 2.1. However, if there is a possibility of binding ZLB at time 1, the threshold for a binding ZLB is $\rho^* > \rho^{**}$. We summarize the analysis in the following:

Proposition 1 *Optimal policy is looser today when the probability of a binding ZLB tomorrow is positive.*

This result has several predecessors; perhaps the closest is Adam and Billi (2007) 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. This result has been correctly interpreted to mean that, if negative shocks lead the economy to be close to the ZLB, the optimal response is to reduce the interest rate aggressively, rather than "sparing dry powder", to reduce the likelihood that the ZLB becomes an effective constraint. However, the same logic applies to liftoff. Following an episode where the ZLB has been a hard constraint on policy, one should not raise rates as if the ZLB constraint was gone forever. Even though the best forecast may be that the economy will recover and exit the ZLB - *i.e.*, in the context of the model, that $E(\rho_1^n) > 0$ – it can be optimal to have zero interest rates today. Note that policy is looser when the probability of falling back in a ZLB situation is high or the potential severity of the ZLB problem is large, *i.e.* $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho$ is a large negative numbers; the economy is less sensitive to interest rates (high σ), and the Phillips curve is steep (high κ).

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

Proposition 2 *Higher uncertainty about the natural real rate tomorrow ρ_1^n leads ceteris paribus to a looser policy today.*

To see this, rewrite the key quantity $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho = E \min(\rho, 0)$. Since the min function is concave, higher uncertainty about ρ leads to lower (*i.e.* more negative) $Ex_1, E\pi_1$ and hence lower i_0.

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

Proposition 3 *Optimal policy is independent of the distribution of the natural rate of interest tomorrow ρ_1^n over values for which the ZLB does not bind,* i.e. *of $\{f_\rho(\rho)\}_{\rho \geq 0}$; only $\{f_\rho(\rho)\}_{\rho<0}$ is relevant, and only through the sufficient statistic $\int_{-\infty}^{0} \rho f_\rho(\rho) d\rho$.*

We now present some simple examples that illustrate our results so far.

Example 1 *Suppose ρ_1^n can take two values: $\rho_1^n = \overline{\rho} + h > 0$, with probability $1/2$, or $\rho_1^n = \overline{\rho} - h$, with probability $1/2$. Assume that $\overline{\rho} + h > 0$ and $\overline{\rho} - h < 0$. Then $Ex_1 = \frac{1}{\sigma}\frac{1}{2}(\overline{\rho} - h) < 0$, and an increase in h (a mean-preserving spread in ρ_1^n) leads to lower Ex_1, and hence $E\pi_1$ and i_0.*

Example 2 *Suppose ρ can take three values at time 1: $\overline{\rho} + h > 0$, or $\overline{\rho}$, or $\overline{\rho} - h$. Assume the probabilities are $p, 1-2p$, and p respectively. The mean is $\overline{\rho}$. Also assume that $\overline{\rho} > 0$ and $\overline{\rho} - h < 0$ so that the ZLB binds only in the third state. Then $Ex_1 = \frac{1}{\sigma}p(\overline{\rho} - h) < 0$ and an increase in p or h will lead policy to be looser.*

Example 3 *Suppose ρ can take three values at time 1: either $\rho + x > 0$, or $\rho > 0$, or $\rho - h < 0$, with probabilities $q, 1-q-p$, and p. Optimal policy is independent of x and q and depends only on $p(\rho - h)$. Both the mean of future natural rate, and the upside risk, are not relevant for policy.*

2.2.3 Discussion of the result and policy implications

First, note that while we deliberately focused on a very simple example, it holds under much more general conditions. For instance, the exact same results hold still if $\{\rho_t^n\}_{t\geq 2}$ follows an arbitrary stochastic process as long as it is positive. We discuss below how the same results hold when the uncertainty is over cost-push shocks rather than the natural rate. 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.

There are two obvious limitations to this result. First, it requires that the central bank is able to offset demand shocks (outside the ZLB), and that there is no cost to doing so (*i.e.* the "divine coincidence" holds). However, while the divine coincidence greatly simplifies the analysis, we do not think it is crucial for our results. Second, this assumes that there is no cost to raising rates quickly if needed. That is, our welfare criterion does not give any value to interest rate smoothing. The policy recommendation to reduce the interest rate when there is more uncertainty naturally implies (in the model) that the rate will rise on average faster over time once the economy recovers.

One reading of these results is a policy implication, namely that interest rates should stay lower than normal, that is lift-off should be delayed, but the central bank should be prepared to raise rates quickly if the economy actually picks up strongly. In other words, policy should be *very* state-contingent in the current situation. This is qualitatively very similar to the policy outlined in the introduction.

Finally, in the current situation, the model implies optimal policy involves having a larger output gap than otherwise to offset the deflationary effects of a possibly binding future ZLB. The recent decline in inflation and in measures of inflation compensation might be explained by expectations of the ZLB binding in the future.

2.2.4 Extensions to uncertainty about potential output and cost push shocks

What does it mean that there is uncertainty about the real natural rate tomorrow ρ_1^n? To understand this better, recall that

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

Uncertainty about ρ_1^n may arise either because it is difficult to know in advance the realizations of some demand components (such as foreign demand, or the components affected heavily by fiscal policy), or because we are uncertain today regarding the future growth of potential output. For instance, if we are unsure whether potential growth will pick up or fall, we face uncertainty about the natural real rate of interest.

A closely related but conceptually different question is, what if we are uncertain about the current level of potential output z_0. For instance, there is currently a large debate among economists on the amount of slack in the economy. We plan to add a result or at least some calculations when there is uncertainty about the level of potential (in the spirit of Ehrmann and Smets (2003)).

In the appendix we consider the case of optimal policy with uncertainty about cost-push inflation. We show that in this case as well optimal policy is looser if there is a chance of a binding ZLB in the future due to a low cost push shock. Another implication of this case is that the risk that inflation picks up due to a high cost push shock does not affect policy today. If such a shock were to occur tomorrow, it will lead to some inflation; however, there is nothing that policy *today* can do about it.

2.3 The buffer stock channel

This section presents a different rationale for loose policy when the economy is close to the ZLB. The expectations channel discussed in the previous section arises because forward-looking agents expect a chance of low inflation and output tomorrow if the ZLB binds, leading to less spending and less inflation today. In contrast, the buffer stock channel does not rely on forward-looking behavior on the part of agents, but rather on the (somewhat opposite) view that the economy has some inherent "momentum". Suppose that output or inflation have an inherent 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. To see this, note that in the standard forward-looking New Keynesian model, the ZLB hits if the following inequality binds:

$$x_t \leq E_t x_{t+1} + \frac{\rho_t^n + E_t \pi_{t+1}}{\sigma},$$

and if output or inflation have some inherent momentum, then $E_t x_{t+1}$ depends on past values of x_t, including x_{t-1}; and similarly for $E_t \pi_{t+1}$. Building up output or inflation thus makes it less likely that this condition will bind. Of course, this argument does not hold in the simple, purely forward-looking model that we discussed in the previous section; it requires that inflation and/or the output gap be at least partly backward-looking. There is ample evidence that such backward-looking terms are important.[12]

The preceding argument does not guarantee that it is optimal to increase output or inflation, for at least two reasons. First and most obviously, the benefit of a high inflation or output today in the event that a ZLB event arises tomorrow must of course be weighted against the costs of excess output and inflation (*i.e.* the cost of "overheating" the economy) today, and tomorrow's cost to bring down the output gap or inflation if the economy turns out not to hit the ZLB constraint. As we might guess, these costs will be second-order if we start close to target. Second, the argument for loosening policy is actually more complicated, because the momentum affects not just the ZLB constraint, but also the Phillips curve; and even the effect on the constraint can be subtle because both expected output and expected inflation show up there. In short, we need to demonstrate that this simple intuition is valid, at least in some cases.

We have so far worked out three models in which this argument holds, at least for a range of parameters. Here we present the simplest case of a completely backward-looking model. In the appendix we study two more models. One where the IS curve is backward-looking,

[12] Indeed, most medium-scale DSGE models incorporate both forward-looking and backward-looking terms, in the form

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

and the presence of the lagged terms can be microfounded by appealing to habits (for the IS curve) and to price indexation (for the Phillips curve). These backward-looking terms are found to be empirically important both in the estimated DSGE models and in the single-equation estimations.

but the Phillips curve is forward-looking. The other with a backward-looking Phillips curve and forward-looking IS curve.

We consider a traditional backward-looking model, where inflation and the output gap are affected by lagged values of inflation and output gap, rather than by future values. The model is:[13]

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

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

We calculate optimal policy when there is uncertainty about the natural real rate of interest only. We start at time 2. The economy does not experience any more shocks, so $\rho_t^n = \rho$, but it starts with an initial "inflation momentum" π_1 and an "output gap momentum" x_1 (both of which may be positive or negative) . The optimal policy under discretion involves reducing these gaps back to zero. The output gap term can be easily adjusted by changing the interest rate, provided we do not hit the ZLB at time 2, *i.e.* that $\rho_2^n = \rho$ is large enough, an assumption we will maintain.[14] The inflation gap requires an output gap in the opposite direction. Given the quadratic loss, it is optimal to smooth this adjustment over time, so the economy will converge back to its steady-state slowly. The details of this adjustment process after time 2 are not terribly important for our analysis; what is important is that the overall loss of starting from period 2 with a lagged inflation π_1 and output gap x_1 turns out to be a quadratic function of π_1 only, so that we can write it $\frac{V}{2}\pi_1^2$ for some V, where V is a number that depends on λ, κ, and β and is calculated in appendix.

Turn now to optimal policy at time 1. For a given ρ_1^n realization, and given last period's output gap x_0 and inflation π_0, the problem can be written as:

$$V(x_0, \pi_0, \rho_1^n) = \min_{x_1} \frac{1}{2}\left(\pi_1^2 + \lambda x_1^2\right) + \beta \frac{V}{2}\pi_1^2,$$

[13]One may wonder why π_{t-1} rather than π_t shows up in the second equation. However, given the Phillips curve equation, the model is equivalent to $x_t = \widehat{\delta} x_{t-1} - \frac{1}{\sigma}\left(i_t - \rho_t^n - \pi_t\right)$, where $\widehat{\delta} = \delta - \kappa/\sigma$.

[14]Relaxing it would only strengthen our results.

$$
\begin{aligned}
s.t. \quad &: \\
\pi_1 &= \pi_0 + \kappa x_1, \\
x_1 &\leq \delta x_0 + \frac{\pi_0}{\sigma} + \frac{\rho_1^n}{\sigma},
\end{aligned}
$$

where the main change compared to the previous analysis is that the policymaker now anticipates the cost of having inflation π_1 tomorrow, and conversely his choices are affected by yesterday's values x_0 and π_0. Depending on the value of ρ_1^n, two cases can arise. Either ρ_1^n is high enough that the ZLB does not bind, leading to the solution

$$
\begin{aligned}
x_1 &= -\frac{(1+\beta V)\kappa}{(1+\beta V)\kappa^2+\lambda}\pi_0, \\
\pi_1 &= \frac{\lambda}{(1+\beta V)\kappa^2+\lambda}\pi_0.
\end{aligned}
$$

This solution reflects the standard trade-off between output and inflation, except that the weight on inflation $1+\beta V$ now takes into account the cost of having inflation away from target tomorrow. The corresponding loss is

$$
V(x_0, \pi_0, \rho_1^n) = \frac{V}{2}\pi_0^2,
$$

since in this case the problem is actually the same as the one faced at time 2. The alternative case is that the ZLB binds, if ρ_1^n is low enough, in which case

$$
\begin{aligned}
x_1 &= \delta x_0 + \frac{\pi_0}{\sigma} + \frac{\rho_1^n}{\sigma}, \\
\pi_1 &= \kappa\delta x_0 + \pi_0\left(1+\frac{\kappa}{\sigma}\right) + \kappa\frac{\rho_1^n}{\sigma},
\end{aligned}
$$

and the loss is

$$
\begin{aligned}
V(x_0, \pi_0, \rho_1^n) &= \frac{1}{2}\pi_1^2 + \frac{\lambda}{2}x_1^2 + \frac{1}{2}\beta V \pi_1^2, \\
&= \frac{1+\beta V}{2}\left(\kappa\delta x_0 + \pi_0\left(1+\frac{\kappa}{\sigma}\right) + \kappa\frac{\rho_1^n}{\sigma}\right)^2 + \frac{\lambda}{2}\left(\delta x_0 + \frac{\pi_0}{\sigma} + \frac{\rho_1^n}{\sigma}\right)^2,
\end{aligned}
$$

which is greater than $\frac{V}{2}\pi_0^2$ since it is the solution to a constrained problem.

The threshold value for ρ_1^n can be calculated as:

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

which is decreasing in both x_0 and π_0. The probability of hitting the ZLB is

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

Hence, in contrast to the example of the previous section where the probability of hitting the ZLB constraint tomorrow was exogenous, it is now endogenous and can be influenced by policy today. Specifically, a higher output gap or inflation today will reduce the likelihood of hitting the ZLB tomorrow. The expected loss from time 1 on is also a function of the output gap and inflation:

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

This expressions reveals that the initial conditions x_0 and π_0 matter by shifting (i) the payoff from continuation in the non-ZLB states, $\frac{V}{2}\pi_0^2$, (ii) the payoff in the case where the ZLB binds (the term under the second integral), (iii) the relative likelihood of ZLB and non-ZLB states (through $\rho_1^*(x_0, \pi_0)$). Since the loss function is continuous in ρ (even at $\rho_1^*(x_0, \pi_0)$), this last effect is irrelevant for welfare at the margin.[15]

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

[15] Hence, a main goal of optimal policy is to reduce the severity of the loss in the states in which the ZLB binds; reducing the likelihood of hitting the ZLB, while possible, has small benefits.

expected loss tomorrow:

$$\begin{aligned} & \min_{x_0} \frac{1}{2}\left(\pi_0^2+\lambda x_0^2\right)+\beta L(x_0,\pi_0), \\ s.t. \quad & : \\ x_0 \quad & \leq \quad \delta x_{-1}+\frac{\rho_0^n+\pi_{-1}}{\sigma}, \\ \pi_0 \quad & = \quad \pi_{-1}+\kappa x_0. \end{aligned}$$

To simplify the exposition, we now assume that the distribution of ρ_1^n is such that there is only one negative value of ρ, call it $\underline{\rho}$; this arises with probability p; and the solution will be such that this is the only value of ρ for which the ZLB binds. As a result, we can rewrite the expected loss L as

$$L(x_0,\pi_0)=\frac{V}{2}\pi_0^2+p\left(\frac{1+\beta V}{2}\left(\kappa\delta x_0+\pi_0\left(1+\frac{\kappa}{\sigma}\right)+\kappa\frac{\underline{\rho}}{\sigma}\right)^2+\frac{\lambda}{2}\left(\delta x_0+\frac{\pi_0}{\sigma}+\frac{\underline{\rho}}{\sigma}\right)^2-\frac{V}{2}\pi_0^2\right).$$

We can now prove the following:

Proposition 4 *Consider a purely backward-looking model, and suppose the initial lagged inflation $\pi_{-1}=0$. Suppose that there is one value of the natural real rate, ρ_1^n, which causes a zero-lower-bound at time* 1*; call $\underline{\rho}$ this value and denote p its probability. Then, optimal policy is looser when p is higher, or when the severity of the ZLB $\underline{\rho}$ becomes bigger. Specifically, x_0 and π_0 are increasing in p (at least for small p), and decreasing in $\underline{\rho}$, while the optimal interest rate i_0 is decreasing in p (at least for small p) and increasing in $\underline{\rho}$.*

The intuition is simple. For $\pi_{-1}=0$ and $p=0$, the optimal solution is $x_0=\pi_0=0$. As p starts to increase, the cost from smoothing inflation are second-order relative to the benefits in the event that the ZLB binds.

Note that as in the previous section, this result does not rely on assumptions about the mean of ρ_1^n or the distribution of its values when the ZLB does not bind - only the probability and severity of the ZLB problem affect the solution. In the appendix we demonstrate that a result similar to the one here holds even in a model that is partly forward-looking.

2.4 Illustrative simulations

We now demonstrate how uncertainty affects optimal liftoff by simulating the different cases. Currently we have worked out one example involving the forward-looking model, and hence is illustrates only the expectations channel. We will add some examples to illustrate the backward-looking channel and perhaps calculate optimal policy in a mixed model with both backward- and forward-looking elements. While a full quantification of the the two channels is beyond the scope of this paper, these examples will hopefully illustrate that the two mechanisms can have substantial effects.

Suppose that ρ_t^n was negative for $t = -10, ..., -1$, and is finally positive at time 0, with value ρ_0^n. It is forecast to rise and return to its steady-state value, $\overline{\rho}$. However, there is uncertainty about its values at time 1 through $T-1$:

$$\begin{aligned} \rho_T^n &= \overline{\rho} \text{ with probability } 1, \\ \rho_t^n &= \rho_0^n + \left(\overline{\rho} - \rho_0^n\right)\frac{t}{T} + \sigma_\varepsilon \varepsilon_t, \end{aligned}$$

where ε_t is independent over time and distributed according to a truncated normal distribution $N(0,1)$. The truncation is helpful to make sure that, after time T at least, there is no risk of falling back to the ZLB. This allows an easy backward calculation of the optimal policy.

Optimal policy under discretion can be easily calculated in this model. For $t \geq T$, $x_t = \pi_t = 0$. For $t < T$, the optimal policy is given by $i_t = \max\left(\rho_t^n + E_t\pi_{t+1} + \sigma E_t x_{t+1}, 0\right)$. If this interest rate is positive, it yields $x_t = \pi_t = 0$. If not, we obtain $x_t = \frac{\rho_t^n}{\sigma} + E_t x_{t+1} + \frac{E_t\pi_{t+1}}{\sigma}$ and $\pi_t = \beta E_t \pi_{t+1} + \kappa x_t$. The conditional expectations $E_t\pi_{t+1}$ and $E_t x_{t+1}$ are calculated by backward induction. Define $a_t = E_t x_{t+1}$ and $b_t = E_t \pi_{t+1}$. We have $a_{T-1} = b_{T-1} = 0$. Next, define the threshold for the ZLB to bind as $\rho_t^* = -E_t\pi_{t+1} - \sigma E_t x_{t+1} = -b_t - \sigma a_t$. Then,

$$\begin{aligned} a_{t-1} &= E_{t-1}\left(1_{\rho \leq \rho_t^*}\left(\frac{\rho_t^n}{\sigma} + a_t + \frac{b_t}{\sigma}\right)\right), \\ b_{t-1} &= E_{t-1}\left(\kappa x_t + \beta E_t \pi_{t+1}\right) = \kappa a_{t-1} + \beta b_t, \end{aligned}$$

Figure 1: Optimal Lift-off in the Forward Looking Model

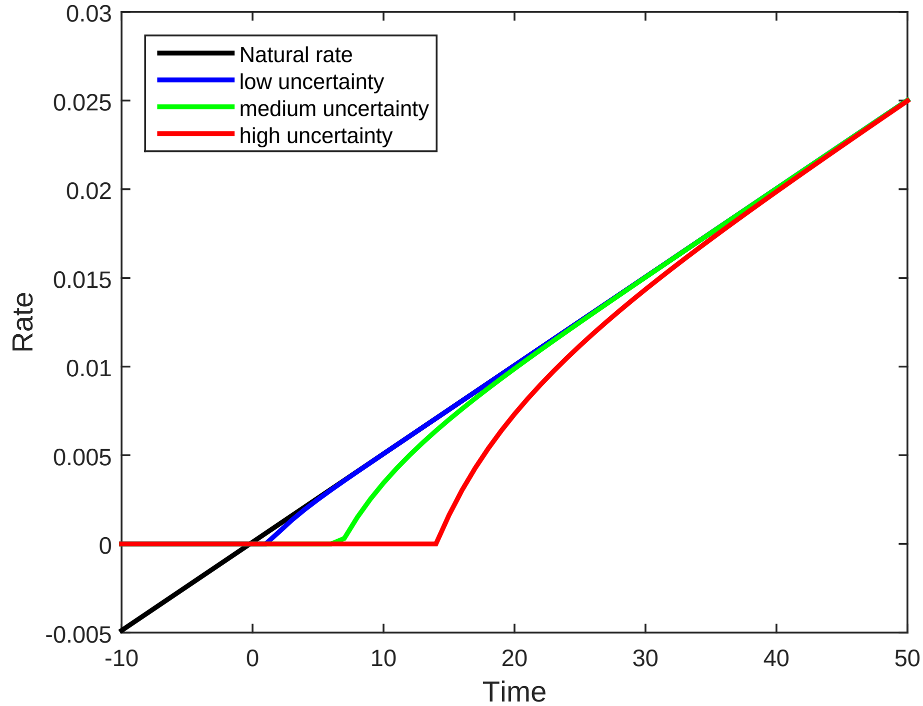

so that knowledge of the distribution of $\{\rho_t^n\}$ and a_t, b_t allows calculating a_{t-1}, b_{t-1}.

Figure 1 displays the path of the expected natural rate $E\rho_t^n$ together with the solution i_t, calculated for different values of σ_ε.[16] If there is no uncertainty, the optimal policy under discretion is to set $i_t = \rho_t^n$. With uncertainty, we see a bias towards lower rates and delayed liftoff.

2.5 Review of the role of uncertainty away from the ZLB

Here we review the literature that provides rationales for a risk-management approach to monetary policy away from the ZLB. There is a long tradition among policymakers to emphasize the importance of uncertainty in their decisions. For instance, Greenspan (2004) argues that "(t)he Federal Reserve's experiences over the past two decades make it clear that uncertainty is not just a pervasive feature of the monetary policy landscape; it is the

[16]We report the interest rate path that arises for $\varepsilon_t = 0$, *i.e.* the shocks take their mean value.

defining characteristic of that landscape." Most of the literature abstracts from the ZLB constraint, which is critical to our own analysis. Hence, this literature is complementary to our work. Since most of our empirical work takes place in periods when the ZLB was not perceived as a likely constraint on monetary policy, this literature is important to explain our findings.

2.5.1 Non-quadratic loss function

Perhaps the simplest possible deviation from the canonical framework that generates a rationale for risk-management is to relax the assumption that the loss function is quadratic. While the quadratic is justified as a local approximation to welfare (see Woodford (2003)), it might not be the correct normative objective function if the shocks are large enough that the local approximation deteriorates. Moreover, the quadratic loss function relies on some simplifying assumptions which may not always hold in reality. To take just one example, it is assumed that there are subsidies that offset the average monopoly power of price-setters. More fundamentally, even if the quadratic loss function is the correct normative approach, it might not provide the best approximation to the way the FOMC behaves.

Motivated by these considerations, several authors have introduced and estimated loss functions that feature asymmetries, e.g. Surico (Surico, 2007), Kilian and Manganelli (2008) Dolado, María-Dolores, and Ruge-Murcia (2004). For instance, the policy-maker may give a higher weight to negative output gap deviations from target than to positive output gap deviations. For some special functional forms, one can show that such a loss function leads to an optimal policy rule that resembles the Taylor rule, but rather than being a linear function of output gap and inflation, it is a nonlinear function. These papers focus on these nonlinearities rather than the implication that uncertainty will affect policy – a separate implication of the non-quadratic loss function.

2.5.2 Non-linear economies

Another possible approach is to relax the assumed linearity of the model. There are different ways in which a nonlinear model might generate a response to uncertainty. One possibility that has been studied empirically and theoretically in the literature is that the Phillips curve may be convex, rather than linear, e.g. Laxton, Rose, and Tambakis (1999), Dolado, María-Dolores, and Naveira (2005). That is, inflation might increase fairly quickly as the output gap becomes positive. This would require policy to respond in a nonlinear fashion to cost-push shocks. Inflation would likely be a convex function of the cost-push shock – a high cost-push shock would trigger a fairly large inflation since it would be too costly to reduce it with a negative output gap given the convexity of the Phillips curve. This nonlinear inflation function would imply that expected inflation is higher ex-ante (given Jensen's inequality) and thus lead to a tighter policy.

An alternative source of non-linearity is expectation formation. During some historical episodes, the FOMC appeared to be concerned with the possibility that inflation expectations might "take off" if inflation remains significantly above target for a while. This would naturally lead to a policy that minimizes this risk. Perhaps the simplest approach to model this story is to assume that expectation formation has a threshold rule, where expectations are not adjusted as long as $\pi_t \leq \overline{\pi}$ but they become adaptive if $\pi_t > \overline{\pi}$. This asymmetry in inflation expectations would likely create a tightening bias in policy – through the same mechanism (though the opposite direction) as the ZLB asymmetry. We are not aware of research providing empirical evidence in favor of this mechanism.

The effect of uncertainty can also be to create an additional shock to the economy. 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. Optimal monetary policy would naturally need to react to this same shock, either to accommodate it or to try and offset some of its effects. One channel is that higher uncertainty may lead to precautionary savings which depresses demand and hence leads to a lower natural rate *today*, see Fernández-Villaverde, Guerró-Quintana, Kuester, and Rubio-Ramírez (2012) or

Born and Pfeifer (2014).

Finally, we note that some recent papers have studied fully non-linear models at the ZLB, including Nakata (2013), Plante, Richter, and Throckmorton (2013), Basu and Bundick (2013) among others. In these models, the inability of the central bank to respond to shocks leads endogenously to higher uncertainty, which may feed back to agents' decisions in the form of higher precautionary savings. These papers focus on this positive implication rather than on the normative implications for policy.

2.5.3 Parameter uncertainty and model uncertainty

A large literature on optimal monetary policy discusses how the policymaker should react if he is uncertain about the true model of the economy. Brainard (1967) derived the important result that parameter uncertainty would lead to additional caution and less policy response to deviations from target, a principle that is often called "gradualism." This principle has had considerable influence on policymakers, for instance Blinder (1998) or Williams (2013). However, it is not uniformly valid as discussed by Barlevy (2011). More generally, a recent and fast-growing literature incorporates concern about model miss-specification into optimal monetary policy analysis, sometimes along the lines suggested by Hansen and Sargent (2008). This model uncertainty may under some conditions lead policy to be more aggressive, but as Barlevy (2011) explains, it might also lead to more caution. Hence, the effect of parameter and model uncertainty are themselves uncertain. Meyer, Swanson, and Wieland (2001) and Swanson (2004) construct examples where the policymaker learns about the natural rate, but the learning is nonlinear because of deviations from the linear Gaussian Kalman filter.[17] In this case, the optimal policy will deviate from the certainty equivalence principle.

[17] In the Meyer et al. (2001), this is justified by the idea that the policymaker is certain that the natural rate lies between 4 and 6%, and has a uniform flat prior in that region.

2.5.4 Reputational costs

Finally, a policymaker must take into account the effect that shocks might have on her reputation; in particular, policymakers may face large costs of reversing a decision. Empirically, it is well known that central banks tend to go through "tightening" and "easing" cycles, *i.e.* there is substantial persistence in the short-term interest rate. One reason why policymakers might be reluctant to reverse course is that it would damage their reputation, perhaps because the public would revise its confidence in the central bank. With high uncertainty, this reputation element would lead to more caution. In the case of liftoff, it argues for more delay to avoid having to revert back to zero.

3 Historical Precedents for Risk Management

The FOMC's historical policy record provides many examples of how risk management considerations have influenced monetary policy decisions. FOMC minutes and other Federal Reserve communications reveal a number of episodes when the Committee appeared to set the stance of policy with more than just the point forecasts for output and inflation in mind. At times, the FOMC indicated that it took 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 the full force of early policy moves had yet shown through to economic activity and inflation. The policy record also indicates several instances when the Committee said its policy stance was taken in part as insurance against undesirable outcomes; during these times, the FOMC also usually noted reasons why the potential costs of a policy overreaction likely were modest as compared with the scenario it was insuring against.

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

Of course, not all risk management references involve reactions to uncertainty or insurance-based rationales for policy. For example, at times the FOMC faced conflicting policy prescriptions for achieving its dual mandate policy goals for output and inflation. Here, the Committee generally hoped to set policy to better align the risks to the projected deviations from the two targets—an interesting balancing act, though not necessarily a violation of the prescription of certainly equivalence policy.

The remainder of this section describes the two episodes we find particularly revealing about the use of risk management in setting rates. We then consider two approaches to quantifying the role of risk management in policy decision-making as it is described in the FOMC minutes for each meeting from 1992 to 2008. The start date is predicated on the fact that FOMC minutes prior to 1992 provide little information about the rationale for policy decisions. Indeed the number of sentences in the minutes prior to 1992 are substantially less than afterwards.

3.1 1997–1998

1997 was a good year for the U.S. economy: real GDP increased 3-3/4 percent, the unemployment rate fell to 4.7 percent—about 3/4 percentage point below the Board of Governors staff's estimate of the natural rate–and core CPI inflation was 2-1/4 percent.[18] But with growth solid and labor markets tight, the FOMC clearly was concerned about a buildup in inflationary pressures. As noted in the Federal Reserve's February 1998 Monetary Policy Report:

> The circumstances that prevailed through most of 1997 required that the Federal Reserve remain especially attentive to the risk of a pickup in inflation. Labor markets were already tight when the year began, and nominal wages had started

[18]The GDP figure refers to the BEA's third estimate for the year released in March 1999.

> to rise faster than previously. Persistent strength in demand over the year led to economic growth in excess of the expansion of the economy's potential, intensifying the pressures on labor supplies.

Indeed, over much of the period between early 1997 and mid-1998, the FOMC directive maintained a bias indicating that it was more likely to raise rates to battle inflationary pressures than it was to lower them. Nonetheless, the FOMC left the federal 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 economic growth to moderate to a more sustainable pace as well as the fact that actual inflation had remained contained despite tight labor market conditions.[19] But, in addition, on numerous occasions heightened uncertainty over the outlook for growth and inflation apparently reinforced the decision to refrain from raising rates. The following quote from the July FOMC 1997 minutes is an example:

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

Furthermore, the Committee did not see high costs to "waiting and seeing." They thought any increase in inflation would be slow, and that if needed a limited tightening on top of the

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

current 5.5 percent funds would be sufficient to reign in any emerging price pressures. 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, at least on the basis of some measures such as the level of real short-term interest rates. 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 in addition to baseline projections, uncertainty and associated risk management considerations supported the Committee leaving policy on hold.

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

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

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

> (September) ... such an action was desirable to cushion the likely adverse consequences on future domestic economic activity of the global financial turmoil that had weakened foreign economies and of the tighter conditions in financial markets in the United States that had resulted in part from that turmoil. At a time of abnormally high volatility and very substantial uncertainty, it was impossible to predict how financial conditions in the United States would evolve.... In any event, an easing policy action at this point could provide added insurance against the risk of a further worsening in financial conditions and a related curtailment in the availability of credit to many borrowers.
>
> (November) The members were more concerned, however, about the risks stemming from the still sensitive state of financial markets, and in that regard many believed that a prompt policy easing would help to ensure against a resurgence of severe financial strains. A further easing move would complete the policy adjustment to the changed economic and financial climate that had emerged since midsummer and would provide some insurance against any unexpectedly severe weakening of the expansion.

While these references to insurance are clear, the case also can be made that these policy moves were mainly responses to changes in the probability distributions on the outlooks for growth and inflation. Over this time the policy prescriptions to address the risks to the FOMC's dual mandate policy goals were in conflict—risks to achieving the inflation mandate called for higher interest rates while risks to achieving the maximum employment mandate called for lower rates.[20] As the above quote from Chairman Greenspan indicated, in 1997 the Committee thought that a 5-1/2 percent funds rate kept these risks in balance. But as the odds of economic weakness increased, the Committee cut rates to bring the risks to the outlook back into balance. As Chairman Greenspan indicated in his February 1999 Monetary Policy Testimony:

> To cushion the domestic economy from the impact of the increasing weakness in foreign economies and the less accommodative conditions in U.S. financial markets, the FOMC, beginning in late September, undertook three policy easings. By mid-November, the FOMC had reduced the federal funds rate from 5-1/2 percent to 4-3/4 percent. These actions were taken to rebalance the risks to the outlook, and, in the event, the markets have recovered appreciably.

So were the late 1998 rate moves a balancing of forecast probabilities, insurance, or some of both? There is no easy answer. This motivates our econometric work in Section 4 that seeks to disentangle the normal response of policy to expected outcomes from its response to uncertainty about those outcomes.

In the end, the economy weathered the fallout from the Russian default well. In June 1999, the staff forecast projected the unemployment rate to end the year at 4.1 percent and that core CPI inflation would rise to 2.5 percent by 2000.[21] Against this backdrop,

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

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

the FOMC decided to increase the funds rate to 5 percent. In the event, the staff forecast underestimated the strength of the economy and underlying inflationary pressures, and the FOMC ended up executing a series of rate hikes that eventually brought the funds rate up to 6.5 percent by May of 2000.

3.2 2000–2001

At the time of the June 2000 FOMC meeting, the unemployment rate stood at 4 percent and core PCE inflation, which the Committee was now using as its main measure of consumer price inflation, was running at about 1-3/4 percent, up from 1-1/2 percent in 1999. The staff forecast growth would moderate to a rate near or a little below potential but that unemployment would remain near its current level and that inflation would rise to 2.3 percent in 2001—and this forecast was predicated on another 75 basis points tightening that would bring the funds rate to 7-1/4 percent by the end of 2000.

Despite this outlook, the FOMC decided to leave rates unchanged. What drove this pause? It seems likely that more than the forecast of moderating growth was involved, and that risk management also was an important consideration.

In particular, the FOMC appeared to put further rates hikes on hold in part 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, but the degree of the slowing was not clear. Second, rates had risen substantially over the past year, and given the lags from policy changes to economic activity, it was unlikely that the full effects of the hikes had yet been felt. Given the relatively high level of the funds rate and the slowdown in growth that appeared in train, the Committee appeared wary of over tightening. Third, despite the staff forecast, it was noted that the FOMC considered the costs of waiting in terms of inflation risks to be small. Accordingly, they thought it better to put a rate increase on hold and see how the economy developed. The June 2000 minutes contain a good deal of commentary supporting this interpretation:[22]

[22] This was not the first time the Committee had invoked such arguments during the tightening cycle.

> The increasing though still tentative indications of some slowing in aggregate demand, together with the likelihood that the earlier policy tightening actions had not yet exerted their full retarding effects on spending, were key factors in this decision. The uncertainties surrounding the outlook for the economy, notably the extent and duration of the recent moderation in spending and the effects of the appreciable tightening over the past year, including the 1/2 percentage point increase in the intended federal funds rate at the May meeting, reinforced the argument for leaving the stance of policy unchanged at this meeting and weighting incoming data carefully. . . .Members generally saw little risk in deferring any further policy tightening move, particularly since the possibility that underlying inflation would worsen appreciably seemed remote under prevailing circumstances. Among other factors, inflation expectations had been remarkably stable despite rising energy prices, and real interest rates were already relatively elevated.

Moving through the second half of 2000, it became increasingly evident that growth had slowed to a pace somewhat below trend and may in fact have been poised for even more pronounced weakness. Furthermore, inflation was moving up at a slower pace than the staff had projected in June. In response, the Committee held 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 the Committee had anticipated. In a conference call on January 3, 2001, the FOMC cut the funds rate to 6 percent and lowered it again to 5-1/2 percent at the end-of-month FOMC meeting.[23] In justifying the aggressive ease, the Committee stated:

> . . . all the members endorsed a proposal calling for a further easing in reserve conditions consistent with a 50 basis point decrease in the federal funds rate

In October 1999 the FOMC left rates unchanged in part over uncertainty over the economic outlook. And in the February and March 2000 meetings they opted for small 25 basis point cuts because of uncertainty. As stated in the July 2000 Monetary Policy Report to Congress regarding the smaller moves in February and March: "The FOMC considered larger policy moves at its first two meetings of 2000 but concluded that significant uncertainty about the outlook for the expansion of aggregate demand in relation to that of aggregate supply, including the timing and strength of the economy's response to earlier monetary policy tightenings, warranted a more limited policy action."

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

> to a level of 5-1/2 percent. 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 provided by the Committee's policy easing actions would help guard against cumulative weakness in economic activity and would support the positive factors that seemed likely to promote recovery later in the year.... In current circumstances, members saw little inflation risk in such a "front-loaded" easing policy, given the reduced pressures on resources stemming from the sluggish performance of the economy and relatively subdued expectations of inflation.

According to this quote, not only was the actual weakening in activity an important consideration in the policy decision, but uncertainty over the extent of the downturn – and the possibility that it might turn into an outright recession – seemed to spur the Committee to make a large move. The "help guard against cumulative weakness" language could be read as the Committee taking out some additional insurance against the possibility that the weakening activity would snowball into a recession.[24] Furthermore, the Committee thought this insurance was cheap, as the risks of inflation taking off were quite small.

The FOMC steadily brought the funds rate down further over the course of 2001 against the backdrop of a weakening activity, though the economy still 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 1-1/2 percent at an annual rate and rise at just a 1-1/4 percent rate in the first half of 2002. By the end of 2002 the unemployment rate was projected to rise to 6.1 percent and core PCE inflation was projected to be 1-1/2 percent. These forecasts

[24] The Staff forecast made in January 2001 did not quite reach a recession threshold. It predicted that real GDP would fall at a 1/2 percent annual rate in 2001:Q1 but then would rise at a 1.3 percent in 2001:Q2 and a 3.1 percent pace in the second half of the year.

were predicated on the federal funds rate remaining flat at 2-1/4 percent.

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

> Most members, however, favored a 50 basis point reduction in the Committee's target federal funds rate. These members stressed the absence of evidence that the economy was beginning to stabilize and some commented that indications of economic weakness had in fact intensified. Moreover, it was likely in the view of these members that core inflation, which was already modest, would decelerate further. In these circumstances insufficient monetary policy stimulus would risk a more extended contraction of the economy and possibly even downward pressures on prices that could be difficult to counter with the current federal funds rate already quite low. Should the economy display unanticipated strength in the near term, the emerging need for a tightening action would be a highly welcome development that could be readily accommodated in a timely manner to forestall any potential pickup in inflation.

This passage suggests that the large cuts were aimed at preventing the economy from developing self-reinforcing dynamics that could accumulate into a serious recession. Indeed, it notes the costs of such an outcome would be quite high because of limited scope for policy reaction imposed by the ZLB on interest rates. So, the aggressive policy moves could be read at least in part as action by the Committee to take out insurance against downside tail events.

3.3 Quantifying References to Uncertainty and Insurance in FOMC Minutes

Clearly, the minutes contain many references to the Committee noting that uncertain economic conditions influenced their policy decision and times when insurance was cited as a reason to alter the stance of policy one way or the other. The challenge is to quantify these considerations into variables that can be used in empirical work.

In the spirit of the narrative approach pioneered by Romer and Romer (1989, 1994), we built judgmental indicators based on our reading of the minutes. We concentrated on the paragraphs that describe the Committee's rationale for its policy decision, reading these passages for references to when insurance considerations or uncertainty over the economic environment or the efficacy of current or past policy moves appeared closely linked to the FOMC's decision. Other portions of the minutes were excluded from our analysis– for example, the parts that cover staff and participants' views of current and prospective economic and financial developments – in order to better isolate arguments that directly influenced the meeting's policy decision from more general discussions of unusual data or normal forecast uncertainty.

We constructed two separate indicator variables—one for uncertainty (UncIndex) and one for insurance (InsIndex). The uncertainty variable was coded to plus one if we judged that the Committee positioned the funds rate higher than it otherwise would due to uncertainty. We coded a minus one if it appeared that uncertainty caused the FOMC to put rates lower than they otherwise would be. If uncertainty did not appear to be an important factor influencing the policy decision, we coded the indicator as zero. We similarly coded the insurance variable by identifying when the minutes cited insurance against some adverse outcome as an important consideration in the Committee's decision, again with a value of one meaning rates were higher and a value of minus one meaning they were lower than they otherwise would have been.[25] Since these two variables were never coded differently from

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

zero for the same meeting we also consider their sum (UncInsIndex).

Note that we did not attempt to code a variable for risk management per se. The minutes often contain discussions of risks to the Committee's dual mandate goals. But when not accompanied by references to uncertainty or insurance, the risk management language may simply describe policy settings that balance conflicting risks to the outlooks of output and inflation relative to their implicit targets. Such policy moves may just be adjusting the expected losses along output and inflation paths in a balanced fashion, and so need not be associated with a deviation from certainly equivalence policies.

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

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

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

However, we did not code all mentions of uncertainty as a one or minus one. For example, in March 1998—a meeting when the FOMC did not change rates despite some concern over higher inflation—the Committee did refer to uncertainties over the economic outlook and say that it could wait for further developments before tightening. The FOMC had held the funds rate flat at 5.5 percent for about a year, and so was not obviously in the midst of a tightening cycle; the baseline forecast articulated in the policy paragraphs seemed consistent with the current funds rate setting; and the commentary over the need to tighten was in reference to an indefinite point in the future as opposed to the current or subsequent FOMC meeting. So, in our judgment, uncertainty did not appear to be a very important factor holding back a rate increase at this meeting and we coded this date as a zero. Quoting the minutes (again, with our emphasis added):

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

Of course, such judgments always can be debated, and there is no definitive way to judge the accuracy of the decisions. So we also constructed objective measures of how often references to uncertainty or insurance appeared in the policy paragraphs of the minutes. In particular we constructed conditional measures which count the percentage of sentences containing words related to uncertainty or insurance in conjunction with references to economic activity or inflation. The words we used to capture uncertainty are "uncertainty," "uncertain" and "uncertainties." To capture insurance we used "insurance," "ensure," "assurance" and "risk management." The conditioning words for inflation were "inflation," prices," "deflation" and "disinflation." To condition on activity we used "activity" and "growth." [26] We combined the counts for uncertainty and insurance into the two variables UncWords and InsWords. In addition we consider UncInsWords which is the percent of sentences in the policy paragraph of the minutes that mention fall into any of the four classifications.

Figure 2: Uncertainty Word Count and Indicator

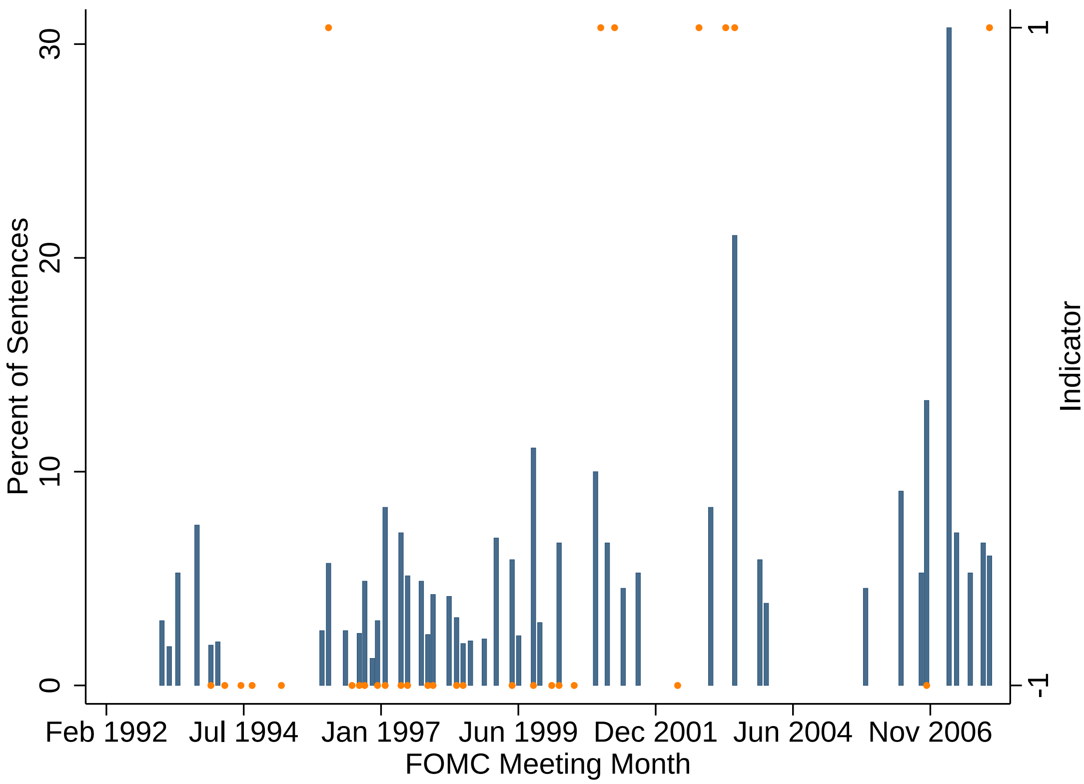

Figures 2 and 3 show plots of these uncertainty and insurance measures. Non-zero values of the indicator variables are indicated by orange circles and the blue bars indicate the word counts. For the word counts we have added together the conditional measures, so for example the insurance word counts reflect mentions of insurance words and at least one of the conditioning words for inflation and activity. Not surprisingly, dealing with uncertainty

[26] After reading the minutes, we realized our word counts failed to account for several commonly used terms, and we will revise these measures in our subsequent draft. In particular for uncertainty we will add "question" and "questions," for inflation we will add "cost" and "costs" and for activity we will add "slack," "resource," "resources," "labor" and "employment."

Figure 3: Insurance Word Count and Indicator

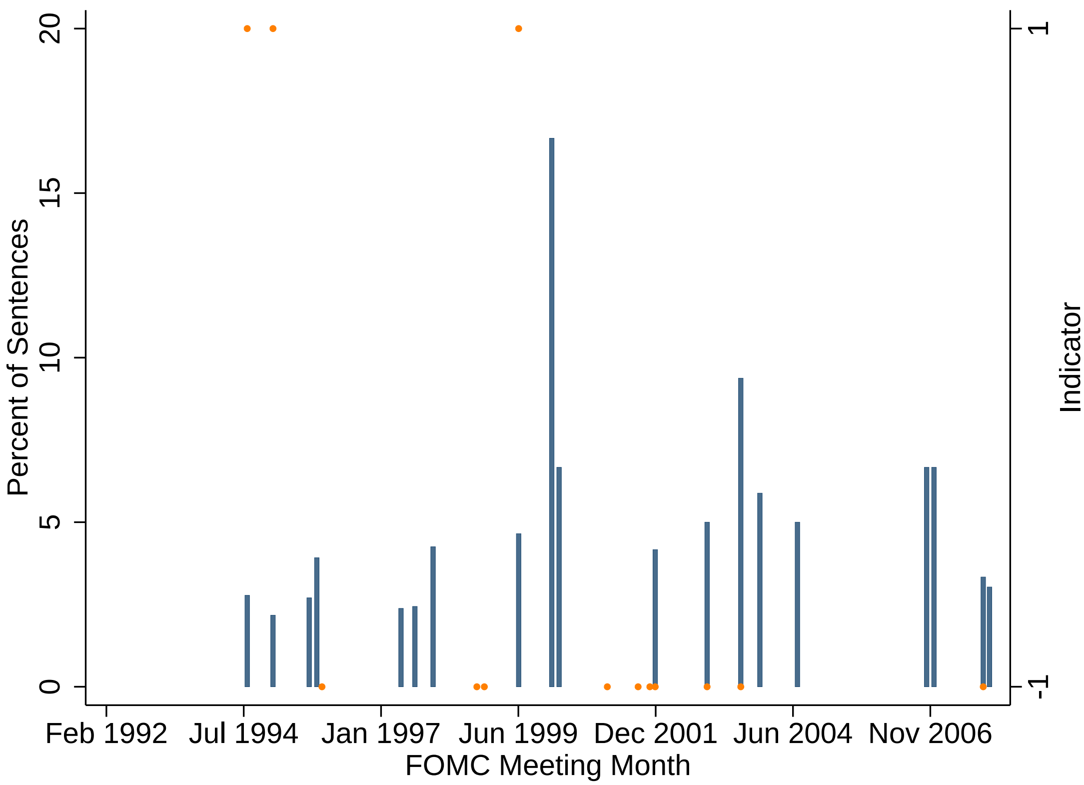

is a regular feature of monetary policy decision making. The uncertainty indicator "turns on" in 31 out of the 132 meetings between 1992 and 2008. Indications that insurance was a factor in shading policy are not as common, but still show up 14 times in the indicator. Most of the time—24 for uncertainty and 11 for insurance—it appears that rates were set lower than otherwise would have been to account for these factors.

The word counts and indicator variables do not line up perfectly. Sometimes the indicator variables are reflected in the words counts but sometimes they are not. There are also meetings where the word counts are positive but we did not judge them to indicate that rates were set differently than otherwise. For example, in March of 2007, our judgmental measure does not code uncertainty as being an important factor putting rates higher or lower than they otherwise would be whereas the word count finds uncertainty referenced in nearly one-third of the sentences in the policy section of the minutes. Incoming data on economic activity were soft, and the Committee was uncertain over the degree to which the economy was weakening. At the same time, they had a good deal of uncertainty on whether their expected decline in inflation – which was running uncomfortably high at the time – actually would materialize. In the end, they only removed the bias in the statement towards further tightening, and did not adjust policy one way or the other in response to the conflicting uncertainties. Hence the judgmental indicator did not code policy being higher of lower than it otherwise would be due to uncertainty.

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

4 Econometric Evidence on Risk Management in Monetary Policy

The previous section delved into the FOMC minutes to uncover instances when the setting of policy was influenced by risk management considerations. While risk management has appeared in the communications of the FOMC it is less clear that it has had an impact quantitatively. For example the FOMC may use risk management arguments to rationalize a particular policy setting but in fact act according to a canonical certainty equivalence policy rule. If this were true then our proposal to incorporate risk management in the current policy environment would be a departure from the prior conduct of the FOMC perhaps making it harder to justify adopting it. The objective of this section is to explore whether risk management has indeed had a material impact on the setting of the Federal funds rate.

Our empirical strategy is straightforward. We formulate and estimate a simple monetary policy reaction function along the lines considered by Clarida, Gali, and Gertler (1998); Clarida et al. (1999); Clarida, Gali, and Gertler (2000) and widely used elsewhere in the empirical literature on monetary policy. This conventional reaction function does not allow a role for anything other than expected values in the setting of the funds rate. To investigate whether other moments of the forecast outlook affect the setting of the funds rate we simple add a variable that proxies for these other moments to the policy reaction function, re-estimate the equation, and determine whether the coefficient on the variable is statistically significant. Statistical significance is interpreted as evidence in favor of our null hypothesis that risk management broadly conceived has in fact been a determining factor in the setting of monetary policy, shading the funds rate decision one way or the other relative it would have been set absent the risk management considerations. This approach comes with drawbacks not least of which is that theory suggests that uncertainty can attenuate or amplify policy responses to incoming information.

We will consider a broad array of variables as proxies for risk management. The analysis of FOMC minutes in the previous section indicates that while risk management has clearly been part of FOMC communications, it is nevertheless challenging to pin down precisely a measure of risk management. Nevertheless as discussed above we have used the minutes to construct two kinds of proxies for risk management and these variables are among those we include in our empirical analysis. While still preliminary (additional work to follow is indicated below), our findings using these proxies show only very weak evidence in favor of our hypothesis that risk management has been a quantitatively important factor in the setting of the funds rate. This fact and the recognition that the nuanced nature of FOMC communications are inherently difficult to quantify leads us to consider several more indirect proxies for risk management. These findings are more favorable to our hypothesis.

The remainder of this section describes our empirical framework in more detail, the proxies for risk we work with, and then reports our estimates of policy rules using the various risk management variables.

4.1 Empirical Strategy

We employ a policy reaction function used frequently in the literature to test the null hypothesis that risk managment has had a material impact on the setting the federal funds rate. Let R_t^* denote the target rate for the nominal federal funds rate in period t. We assume the FOMC uses the following rule for setting its target:

$$R_t^* = R^* + \beta \left(E_t \left[\pi_{t,k}\right] - \pi^*\right) + \gamma E_t \left[x_{t,q}\right] + \xi s_t, \tag{1}$$

where $\pi_{t,k}$ denotes the annualized percent change in the price level between periods t and $t+k$, π^* is the FOMC's target for inlation, $x_{t,q}$ is a measure of the average output gap between period t and $t+q$, s_t is a risk managment proxy,and $\mathcal{E}_t$ is the expectations operator conditional on information available to the FOMC at time t. The coefficients β, γ and ξ are assumed to be fixed over time and summarize the responsiveness of the FOMC to the expected inflation and output gaps and higher order moments of the FOMC's forecast of the economic outlook.

By construction R^* is the desired nominal rate when inflation is at its target and the output gap is closed. Under the assumptions that the output and inflation gaps on average are equal to zero and that the real interest rate is determined by non-monetary factors outside the control of the FOMC in the long run implies

$$R^* = r^* + \pi^*$$

where r^* is the exogenously given long run "equilibrium" real interest rate. So our assumption that the desired nominal rate is constant reflects the underlying assumptions that the long run real interest rate and the inflation targets are both constant. Below we describe how we relax these assumption although doing so has no impact on our findings.

Our empirical implimentation of (1) embodies two additional assumptions. First, we assume that the FOMC has a preference for interest rate smoothing and so does not choose the funds rate to hit its target instantaneously. Second we assume the FOMC does not have perfect control over interest rates. This motivates the following specification for the actual funds rate, R_t

$$R_t = (1 - \bar{\rho})R_t^* + \rho(L)R_{t-1} + \upsilon_t \tag{2}$$

where

$$\rho(L) = \rho_1 + \rho_2 L + \rho_3 L^2$$

and

$$0 \leq \bar{\rho} = \rho_1 + \rho_2 + \rho_3 < 1$$

The variable υ_t is a mean zero and serially independent exogenous interest rate shock and R_t^* is given by (1). Allowing for up to three lags in this partial adjustment formulation simplifies our empirical analysis by ensuring that our estimates of υ_t are serially uncorrelated.

Combining the target model (1) with the partial adjustment equation (2) yields the following equation for the nominal funds rate

$$R_t = b_0 + b_1 E_t\left[\pi_{t,k}\right] + b_2 E_t\left[x_{t,q}\right] + \rho(L) R_{t-1} + b_3 s_t + \upsilon_t. \tag{3}$$

where

$$\begin{aligned} b_0 &= (1-\bar{\rho})\left[r^* + (1-\beta)\pi^*\right] \\ b_1 &= (1-\bar{\rho})\beta \\ b_2 &= (1-\bar{\rho})\gamma \\ b_3 &= (1-\bar{\rho})\xi \end{aligned}$$

We use the publicly available Federal Reserve Board of Governors' staff forecasts of core CPI inflation (in percentage points) and the output gap (percentage point deviations of real GDP from its potential) to measure $\mathcal{E}_t\left[\pi_{t,k}\right]$ and $\mathcal{E}_t\left[x_{t,q}\right]$ with $q = 3$.[27] These estimates are available for each of the eight FOMC meetings a year. In this draft we focus exclusively on results based on the staff estimates for the FOMC meetings closest to the middle of each quarter of a year. In the next draft we will include estimates based on data from all the FOMC meetings for which we have data. We use the average funds rate (in percentage points) over the 30 days after an FOMC meeting to measure R_t. Because our measures of $E_t\left[\pi_{t,k}\right]$ and

[27] These are obtained from the Federal Reserve Bank of Philadelphia public web site.

$E_t [x_{t,q}]$ are based solely on information available before an FOMC it follows that we can obtain consistent estimates of β, γ and ξ by estimating (3) by ordinary least squares, as long as there are sufficient lags in R_t to ensure that the errors υ_t are serially uncorrelated.

We use the Board staff's forecasts in our estimation for two reasons. First, we think doing so leads to a plausible characterization of monetary policy. For example, while these forecasts do not necessarily reflect the views of individual FOMC participants in practise (as reflected in FOMC meeting transcripts and minutes) they typically reflect the views of a plurality. In addition, these forecasts are based on data available at the time monetary policy decisions are made. Second, we do not need instrumental variables to obtain consistent estimates as do Clarida, Gali and Gertler in their papers. Gnabo and Moccero (2014) estimate a version of (3) using Board staff forecasts as well.

We test the null hypothesis that risk management has had a material impact on monetary policy setting by estimating (3) and testing whether ξ is significantly different from zero. This approach comes with drawbacks not least of which is that theory suggests that uncertainty can attenuate or amplify policy responses to incoming information. Finding a significant coefficient will suggest one set of factors are dominant in risk management, while finding insignificant coefficients will be inconclusive.

4.2 Proxies for Risk Management

We consider several other proxies for risk management in addition to those discussed above that we derive from FOMC meeting minutes. Since most of them are only available quarterly we translate the FOMC indicators to quarterly measures by using the indicator value from the meeting closest to the middle of the corresponding quarter.

The additional proxies are divided into two groups: uncertainty and asymmetry. The uncertainty variables are measures of forecast variance while the asymmetry variables measure higher order features of the forecast distribution. In this draft our sample is 1992q1 to 2008q4. The start date corresponds to a change in the way the Survey of Professional Forecasters measures uncertainty that occurs in 1992q1 that makes it difficult to integrate data from early surveys which go back to 1968q4. In principal other measures of uncertainty we use are not constrained by this so it is possible to consider longer sample periods for some cases. The uncertainty and asymmetry variables are all measured in percentage points.

We consider five uncertainty variables, as follows.

1. VIX. This is the Chicago Board Options Exchange Market Volatility Index. It uses the prices of call and put options on the S&P 500 index to measure market participants' expectations of stock market volatility over the next 30 day period.[28] The drawback to using this measure is that it possibly confounds uncertainty due to financial factors that could be unrelated to the outlook for the economy. In addition market volatility over a 30 day period may not be informative about the forecast horizon the FOMC uses to set policy. Nevertheless, the S&P 500 is a broad index of stock prices which should reflect expectations of future earnings over relatively long horizons and which because it covers a diversified portfolio of stocks should reflect market participants' uncertainty about the outlook for the economy over horizons similar to that considered by the FOMC when setting policy (as well as interest rates of course). We use the average of daily values within a quarter.

2. UInf. This is constructed using the methodology described in D'Amico and Orphanides (2008) based on the quarterly Survey of Professional Forecasters (SPF). The SPF asks forecasters to describe mean probabilities that outcomes for GDP deflator inflation for the current and following year fall into pre-specified bins.[29] We follow D'Amico and Orphanides (2008)'s procedure to translate these forecast distributions into forecast distributions for inflation over the next four quarters. We then use their statistical procedure to translate these distributions into a summary measure of the uncertainty of the inflation forecast in each quarter of the survey. Their summary measure is an

[28]The VIX is quoted in percentage points and translates, roughly, to the expected movement (with the assumption of a 68% likelihood) in the S&P 500 index over the next 30-day period, which is then annualized. For example, if the VIX is 15, this represents an expected annualized change, with a 68% probability, of less than 15% over the next 30 days.

[29]Ideally we would use measures of uncertainty in consumer price inflation. Unfortunately bin data for consumer price inflation only becomes available late in our sample period.

estimate of the mean across standard deviations of individual forecasters. It addresses in part a key challenge involved with using the SPF to measure underlying uncertainty. In particular the bins are relatively course with widths of 1 percentage point. This means some forecasters place probability on only one or two bins and the true underlying range of a forecast may be narrower than indicated by the endpoints of the lowest and highest bins an individual places probability on or wider given that the lowest and bins are open ended. Despite these drawbacks the SPF is the only source of information on uncertainty and asymmetry in forecast distributions of individual forecasters.

3. UGDP. This is the measure of uncertainty in real GDP forecasts from the SPF estimated analogously to UInf.

4. DInf. This is simply the standard deviation of the point forecasts of GDP deflator inflation from the SPF.[30] This or similar measures of uncertainty are commonly used in the literature.[31] It is truly a measure of forecaster disagreement rather than the underlying uncertainty, but these concepts are related. For example if there is no underlying uncertainty then there will also be no disagreement. Conversely if there is a lot of underlying uncertainty it is likely that there will be more disagreement about the forecast. Due to outliers in the point forecasts we exclude observations below the 10th and above the 90th percentiles of the point forecast distribution in each quarter. Excluding outliers is common in the literature.

5. UGDP. This is the measure of uncertainty in real GDP growth forecasts estimated using the same methodology as DInf.

We consider five measures of asymmetries in the economic outlook as follows.

1. AInf. This is the difference between the mean over the mean forecast of individual forecasters and the mean over the median forecast of individual forecasters based on

[30]Point forecasts of headline CPI are available for a longer sample and we may use this information in the next draft.

[31]See for example Gnabo and Moccero (2014) and the literature they cite.

the same data underlying UInf. A positive (negative) value for AInf indicates that forecasters put greater weight on upside (downside) outcomes to inflation than downside (upside.) We have explored other measures of skewness including third moments and mean less mode but have found these to be no more informative.

2. AGDP. Same as AInf except for real GDP growth.

3. pfAInf. This is the difference between the mean point forecast and the median point forecast from the SPF, excluding outliers as was done for DInf.

4. pfAGDP: Same as pfAInf except for real GDP growth.

5. FcRev: This is the revision to the Greenbook/Tealbook forecast of the average output gap over the following four quarters from one quarter meeting to the next (using obtained from the Federal Reserve Bank of Philadelphia's public web site.) Finding a significant positive coefficient for this variable would indicate that the FOMC responds more strongly to larger shocks than smaller shocks.

We focus on using these variables and the FOMC-minutes-based variables as measures of s_t in (3). However we have considered other specifications designed to approximate possible non-linearities in the policy rule that could emerge from risk management considerations away from the ZLB, for example Brainard-style attenuation of responses to inflation and output gaps because of uncertainty about the impact of policy on activity and inflation. [32] We do not find any significant effects, although there are results in the literature using different methods that suggest non-linear policy rules perform better empirically than linear rules (see the references in Gnabo and Moccero (2014).) The non-linearities typically considered involve policy responding more aggressively if the inflation or growth outlook cross a threshold. These non-linearities do not speak directly to the hypothesis we are interested in as they

[32] These considerations suggest squared inflation and output gap terms as well as interaction of inflation and output gaps with measures of uncertainty should enter significantly into the policy rule. This can be seen by considering the coefficient on the output gap. Suppose this coefficient is a function of the expected size of the gap and uncertainty about the gap going forward. In this case a linear approximation to the coefficient leads directly to including a squared output gap term and and a term where the output gap multiplies a measure of uncertainty about the output gap.

do not address risk. The closest the literature has come to establishing risk as influencing policy is Gnabo and Moccero (2014). They find that policy responds more aggressively and the degree of inertia in policy is lower in periods of high economic risk as measured by VIX.

Table 1 displays summary statistics for inflation, the output gap and the uncertainty and asymmetry variables. On average the expected output gap has been negative and has varied considerably over the sample period. It is essentially uncorrelated with expected inflation. The VIX and FcRev variables have high variances but the other uncertainty and asymmetry variables do not seem to vary very much. Most of the uncertainty and asymmetry variables have small correlations with expected inflation and expected output gap. One interesting exception is that disagreement in the inflation forecast (DInf) has a large negative correlation with the expected output gap. That is, periods when the outlook for activity is deteriorating often correspond to periods when there is a large amount of disagreement about the outlook for inflation.

Table 1: Summary statistics

					Correlation with	
Variable	Mean	Std. Dev.	Min	Max	Inflation	Output Gap
Inflation	2.51	0.48	1.33	3.63	1.00	
Output Gap	-0.37	1.75	-4.4	3.08	-0.03	1.00
VIX	19.57	7.62	11.02	58.68	-0.39	0.06
UInf	0.62	0.06	0.48	0.75	-0.31	-0.19
UGDP	0.77	0.07	0.61	0.96	-0.36	0.05
DInf	0.21	0.05	0.12	0.32	0.05	-0.58
DGDP	0.23	0.07	0.11	0.51	-0.20	-0.17
AInf	0.02	0.04	-0.08	0.11	0.07	-0.06
AGDP	-0.07	0.05	-0.16	0.10	0.16	-0.37
pfAInf	0.01	0.03	-0.06	0.08	0.15	-0.25
pfAGDP	0.00	0.04	-0.09	0.12	0.06	0.26
FcRev	0.01	0.61	-1.95	1.38	0.25	0.27

Table 2 displays cross-correlations of the uncertainty and asymmetry variables. Several correlations are worth noting. First VIX and DGDP are strongly positively correlated. The correlation of VIX with UGDP is somewhat less positive. Both correlations suggests the VIX is a good indicator of uncertainty about the activity outlook. Interestingly the correlation of

VIX and AInf is strongly negative as it is also for VIX and pfAInf. Therefore when markets perceive a lot of uncertainty in the stock market going forward the inflation outlook is skewed to the downside. The correlation of UGDP and UInf is a very large 0.70, and DGDP and DInf are also positively correlated although much less so at 0.35. However DGDP and UGDP are uncorrelated as are DInf and UInf. Evidently disagreement among forecasters is not the same as the average amount of uncertainty they see. Finally, UGDP and AGDP are strongly negatively correlated, at -0.47. High uncertainty in the GDP outlook occurs when forecasts are skewed to the downside.

Table 2: Cross-correlations of Uncertainty and Asymmetry Variables

Variables	VIX	UInf	UGDP	DInf	DGDP	AInf	AGDP	pfAInf	pfAGDP
UInf	0.07								
UGDP	0.29	0.70							
DInf	-0.01	0.11	-0.12						
DGDP	0.58	0.04	0.10	0.35					
AInf	-0.44	0.17	-0.26	0.11	-0.26				
AGDP	-0.06	-0.28	-0.47	0.35	0.11	0.18			
pfAInf	-0.41	-0.10	-0.36	0.26	-0.10	0.28	0.06		
pfAGDP	0.24	-0.24	0.03	0.01	0.13	-0.32	-0.07	-0.07	
FcRev	-0.40	-0.12	-0.18	-0.05	-0.34	0.23	0.05	0.15	0.06

4.3 Policy Rule Findings

Table 3 shows our policy rule estimates with and without the various FOMC-minutes-based variables. The following two tables show estimates with and without the uncertainty and asymmetry variables and each have the same layout as Table 3. The first columns show the policy rule estimates without any risk management variables and the other columns show the results of estimating the policy rule adding one of the risk management variables at a time with the indicated coefficient estimate corresponding to ξ in (1). In the policy rule without any risk management variables the coefficient on inflation (β) is about 2 and on the output gap (γ) is about 0.8. These estimates are highly significant and are similar to other estimates in the literature. The R^2 (not shown) is 0.97.

Table 3 indicates that none of the coefficients on the risk management variables based on the minutes are significant at conventional levels with standard errors about the same size as the coefficient estimates. The point estimates of the words based measures are essentially zero, but the indicator variables' estimates are more meaningful. For example, our coding of uncertainty and insurance measures suggest that when uncertainty or insurance shades policy it does so by about 25 basis points. By focusing on quarterly observations we have excluded half of the FOMC meetings in our sample. Perhaps by estimating policy rules meeting-by-meeting we will obtain more precise estimates of these effects.[33] Overall these results are not a compelling indication that the words of the FOMC are reflected in their deeds.

Table 3: FOMC Minutes Information in Monetary Policy Rules

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ρ_1	1.19***	1.18***	1.17***	1.16***	1.19***	1.19***	1.19***
	(.08)	(.08)	(.08)	(.08)	(.08)	(.08)	(.08)
ρ_2	-.50***	-.49***	-.47***	-.47***	-.51***	-.50***	-.50***
	(.08)	(.08)	(.09)	(.09)	(.08)	(.08)	(.08)
β	2.01***	2.03***	1.97***	2.00***	2.03***	2.01***	2.02***
	(.17)	(.17)	(.17)	(.17)	(.17)	(.18)	(.17)
γ	.79***	.81***	.80***	.83***	.78***	.80***	.78***
	(.07)	(.08)	(.07)	(.08)	(.07)	(.07)	(.07)
UncIndex		.26					
		(.20)					
InsIndex			.44				
			(.51)				
UncInsIndex				.35			
				(.26)			
UncWords					.03		
					(.02)		
InsWords						-.03	
						(.03)	
UncInsWords							.01
							(.02)

Note: Sample is 1992q1 to 2008q4. ***, ** and * indicate statistical significance at the 1, 5 and 10 percent levels respectively. Standard errors are robust to heteroskedasticity.

[33] This is feasible and will be reported in the next draft.

Table 4: Uncertainty in Baseline Monetary Policy Rules

	(1)	(2)	(3)	(4)	(5)	(6)
ρ_1	1.19***	1.07***	1.19***	1.19***	1.19***	.98***
	(.08)	(.11)	(.08)	(.08)	(.08)	(.11)
ρ_2	-.50***	-.24	-.50***	-.50***	-.50***	-.13
	(.08)	(.21)	(.08)	(.08)	(.08)	(.18)
ρ_3		-.13				-.17**
		(.11)				(.08)
β	2.01***	1.70***	2.10***	2.00***	2.03***	1.85***
	(.17)	(.20)	(.20)	(.21)	(.18)	(.20)
γ	.79***	.80***	.80***	.79***	.75***	.74***
	(.07)	(.07)	(.07)	(.07)	(.08)	(.07)
VIX		-.04**				
		(.02)				
UInf			1.87			
			(2.06)			
UGDP				-.14		
				(1.79)		
DInf					-2.06	
					(2.96)	
DGDP						-5.17***
						(1.99)

Note: Sample is 1992q1 to 2008q4. ***, ** and * indicate statistical significance at the 1, 5 and 10 percent levels respectively. Standard errors are robust to heteroskedasticity.

Estimates associated with the uncertainty variables are displayed in Table 4.[34] These show some evidence that risk management has shaded policy away from a typical policy rule. In particular the coefficients on VIX and DGDP are significant at the 5 and 1 percent levels, respectively. In both cases a one standard deviation increase in the variable (indicating more uncertainty) translates into a lower than typical funds rate of about 25 basis points. None of the other variables' coefficients are statistically significant.

Estimates associated with the asymmetry variables are displayed in Table 5. These

[34]Tables 4 and 5 indicate that some of the specifications include a third lag of the funds rate. Our criteria for selecting lag length was based on the presence or lack thereof of statistically significant serial correlation in the residuals. Our rule of thumb is that we increase lag length until the null hypothesis of no serial correlation up to 4 lags is not rejected at the 5% level of significance using the Durbin h-statistic.

Table 5: Asymmetry in Baseline Monetary Policy Rules

	(1)	(2)	(3)	(4)	(5)	(6)
ρ_1	1.19***	1.06***	1.16***	1.19***	1.19***	1.15***
	(.08)	(.11)	(.08)	(.07)	(.08)	(.12)
ρ_2	-.50***	-.18	-.50***	-.48***	-.50***	-.25
	(.08)	(.20)	(.08)	(.08)	(.08)	(.20)
ρ_3		-.19*				-.17*
		(.10)				(.10)
β	2.01***	1.91***	2.08***	1.88***	2.03***	1.82***
	(.17)	(.19)	(.16)	(.18)	(.18)	(.22)
γ	.79***	.76***	.73***	.87***	.80***	.72***
	(.07)	(.07)	(.07)	(.08)	(.08)	(.08)
AInf		2.52				
		(3.28)				
AGDP			-4.58***			
			(1.69)			
pfAInf				9.00**		
				(4.02)		
pfAGDP					-1.93	
					(3.67)	
FcRev						.45
						(.34)

Note: Sample is 1992q1 to 2008q4. ***, ** and * indicate statistical significance at the 1, 5 and 10 percent levels respectively. Standard errors are robust to heteroskedasticity.

estimates also show some evidence that risk management has shaded policy away from a typical policy rule. Here the coefficients on AGDP and pfAInf are significant at the 1 and 5 percent levels. The negative coefficient on AGDP indicates that when forecasters' distributions of future outcomes are skewed upward on average, that is toward higher growth, policy is shaded downward. This is puzzling. The coefficient on pfAInf is easier to interpret; when individual forecasts of inflation are skewed upward the funds rate is shaded higher. In this case a one standard deviation increase in positive skewness translates into a 25 basis points higher funds rate than otherwise.

5 Limits to Unconventional Policies

A key component of our policy proposal is that unconventional policies at the ZLB are not perfect substitutes for conventional policies away from it. This section will discuss a variety of evidence that unconventional policies are viewed this way. We will make this case primarily by studying speeches by Fed officials, dealer surveys, and Blue Chip special questions. This information will be used to make the case that there are widespread doubts about the efficacy of unconventional policies. Furthermore resistance within the FOMC could impose barriers to their use in the future.

We will also discuss the empirical evidence, in particular work by Krishnamurthy and Vissing-Jorgensen (2011, 2013), that shows the primary channel through which LSAPs have their effect is through signalling a commitment to lower rates for longer.[35] The FOMC has also used explicit forward guidance to signal a commitment to keep rates lower for longer. We will present path factor results extending those in Campbell, Evans, Fisher, and Justiniano (2012) to the current period. These will show that the relationship between changes in expected future funds rates and 10 year Treasures after FOMC meetings has remained stable through the ZLB period. Gilchrist, López-Salido, and Zakrajšek (2014) find that the relationship between Treasuries and corporate bonds has also remained stable. Combined these findings suggest it might be possible use forward guidance or signalling with LSAPs to influence real activity if the economy returns to the ZLB soon after a hasty exit. However, exiting from the ZLB prematurely may extinguish the FOMC's credibility which will make it harder to use signally and explicit forward guidance to influence long term interest rates and activity if we fall back to the ZLB.

If unconventional policies are imperfect substitutes for conventional policies and shocks are symmetric (they need not be homoskedastic) then as discussed in Basu and Bundick (2013) and Plante et al. (2013) New Keynesian theory suggests the distribution of real GDP forecast errors should be symmetric away from the ZLB but skewed at or near the ZLB when

[35] This is not to say that there have not been beneficial effects via reduced term premia, see for example D'Amico, English, López-Salido, and Nelson (2012) and D'Amico and King (2013).

the short term interest rate is the only policy tool available to the central bank. Showing that the distribution of forecast errors fits this pattern would be evidence consistent with the hypothesis of imperfect substitutability between conventional and unconventional policies. This can be assessed over a long sample period using SPF forecasts and the real time data set from the Philadelphia Fed. We will present these findings in the next draft.

6 Conclusion

To be added.

References

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

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

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

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

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

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

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

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012). Macroeconomic effects of federal reserve forward guidance. *Brookings Papers on Economic Activity Spring*, 1–54.

Chow, G. (1973). Effect of uncertainty on optimal control policies. *International Economic Review 14*(3).

Clarida, R., J. Gali, and M. Gertler (1998). Monetary policy rules in practice: Some international evidence. *European Economic Review 42*, 1033–1067.

Clarida, R., J. Gali, and M. Gertler (1999). The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature 37*(4), 1661–1707.

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

D'Amico, S., W. English, D. López-Salido, and D. Nelson (2012). The federal reserve's large-scale asset purchase programs: Rationale and effects. *The Economic Journal 122*, 415–446.

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

D'Amico, S. and A. Orphanides (2008). Uncertainty and disagreement in economic forecasting. Finance and Economic Discussion Series, Federal Reserve Board.

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

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

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

Ehrmann, M. and F. Smets (2003). Uncertain potential output: Implications for monetary policy. *Journal of Economic Dynamics & Control 27*, 1611–1638.

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

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

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

Gilchrist, S., D. López-Salido, and E. Zakrajšek (2014). Monetary policy and real borrowing costs at the zero lower bound. NBER Working Paper No. 20094.

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

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

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

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

Krishnamurthy, A. and A. Vissing-Jorgensen (2011). The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity Fall*, 215–265.

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

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

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

Meyer, L. H., E. T. Swanson, and V. W. Wieland (2001). Nairu uncertainty and nonlinear policy rules. *American Economic Review Papers and Proceedings 91*(2), 226–231.

Nakata, T. (2013). Uncertainty at the zero lower bound. Board of Governors Manuscript.

Plante, M., A. W. Richter, and N. A. Throckmorton (2013). The zero lower bound and endogenous uncertainty. FRB Dallas Manuscript.

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

Romer, C. D. and D. H. Romer (1994). Monetary policy matters. *Journal of Monetary Economics 34*(2), 75–88.

Surico, P. Asymmetric reaction functions for the euro area. *Oxford Review of Economic Policy 19*(1).

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

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

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

Swanson, E. (2004). Signal extraction and non-certainty-equivalence in optimal monetary policy rules. *Macroeconomic Dynamics 8*, 27–50.

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

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

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

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

Appendix

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

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

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

$$
\begin{aligned}
& \min_{x_1} \frac{1}{2}\left(\pi_1^2 + \lambda x_1^2\right), \\
s.t. \quad & : \\
\pi_1 &= \kappa x_1 + u_1, \\
x_1 &\leq \frac{\rho}{\sigma}.
\end{aligned}
$$

with the following solution:

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

$$
\begin{aligned}
x_1 &= -\frac{\kappa u_1}{\lambda + \kappa^2}, \\
\pi_1 &= \frac{\lambda u_1}{\lambda + \kappa^2}.
\end{aligned}
$$

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

$$
\begin{aligned}
x_1 &= \frac{\rho}{\sigma}, \\
\pi_1 &= \kappa\frac{\rho}{\sigma} + u_1.
\end{aligned}
$$

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

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

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

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

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

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

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

$$\begin{aligned} & \min_{x_0} \frac{1}{2} \left(\pi_0^2 + \lambda x_0^2 \right), \\ s.t. \quad : \quad & x_0 \leq \frac{\rho_0^n}{\sigma} + Ex_1 + \frac{E\pi_1}{\sigma}, \\ \pi_0 \quad = \quad & \beta E\pi_1 + \kappa x_0 + u_0. \end{aligned}$$

The solution is the following. Define

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

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

$$\begin{aligned} x_0 &= -\frac{\kappa}{\lambda + \kappa^2} \left(\beta E\pi_1 + u_0 \right), \\ \pi_0 &= \frac{\lambda}{\lambda + \kappa^2} \left(\beta E\pi_1 + u_0 \right), \end{aligned}$$

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

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

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

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

Proposition 5 *Suppose the uncertainty is about cost-push shocks. Then: (1) optimal policy is looser today when the probability of a binding ZLB tomorrow is positive; (2) optimal policy is independent of the distribution of the cost-push shock tomorrow u_1^n over values for which the ZLB does not bind,* i.e. *of $\{f_u(u)\}_{u \geq u^*}$; only $\{f_u(u)\}_{u<u^*}$ is relevant, and only through the sufficient statistics $\int_{-\infty}^{u^*} f_u(u)du$ and $\int_{-\infty}^{u^*} uf_u(u)du$.*

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

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

Example 4 *Suppose that u can take two values, $u = +\Delta$ (with probability $1/2$) and $u = -\Delta$ (with probability $1/2$). If Δ is small, then $P = M = 0$, and hence $E\pi_1 = Ex_1 = 0$, and optimal policy is decided taking into account ρ_0^n and u_0 only. If Δ is large enough, then $P = 1/2$, $M = -\Delta/2$, and $Ex_1 = \frac{\rho}{\sigma}\frac{1}{2} - \frac{\kappa}{\lambda+\kappa^2}\Delta$ (which is negative since $-\Delta < u_1^* = \frac{\rho}{\sigma}\frac{\lambda+\kappa^2}{\kappa}$), and $E\pi_1 = \kappa Ex_1$. A higher Δ then reduces $Ex_1, E\pi_1$ and i_0.*

B Calculation of V in the purely backward-looking model

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

$$\begin{aligned} V(\pi_{-1}, x_{-1}) &= \min_{x,\pi} \frac{1}{2}\left(\pi^2 + \lambda x^2\right) + \beta V(\pi, x), \\ s.t. &: \\ \pi &= \pi_{-1} + \kappa x, \\ x &= \delta x_{-1} - \frac{1}{\sigma}(i - \rho - \pi_{-1}). \end{aligned}$$

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

$$V(\pi_{-1}, x_{-1}) = \frac{v_0}{2}\pi_{-1}^2,$$

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

$$\min_x \frac{1}{2}\left(1 + \beta v_0\right)\left(\pi_{-1} + \kappa x\right)^2 + \frac{1}{2}\lambda x^2$$

The first order condition yields

$$x = -\frac{(1+\beta v_0)\kappa}{(1+\beta v_0)\kappa^2+\lambda}\pi_{-1},$$

leading to

$$\pi = \frac{\lambda}{(1+\beta v_0)\kappa^2+\lambda}\pi_{-1},$$

which verifies our guess of linear rules. To find v_0, plug this back in the minimization problem; we look for v_0 to satisfy, for all π_{-1}, :

$$\frac{v_0}{2}\pi_{-1}^2 = \frac{1}{2}(1+\beta v_0)\pi_{-1}^2\left(\frac{\lambda}{(1+\beta v_0)\kappa^2+\lambda}\right)^2 + \frac{1}{2}\lambda\left(\frac{(1+\beta v_0)\kappa}{(1+\beta v_0)\kappa^2+\lambda}\right)^2\pi_{-1}^2$$

i.e.

$$v_0\left((1+\beta v_0)\kappa^2+\lambda\right)^2 = (1+\beta v_0)\lambda^2 + \lambda(1+\beta v_0)^2\kappa^2,$$

and equation which can be simplified to a simple quadratic equation:

$$\beta\kappa^2 v_0^2 + v_0\left(\kappa^2+\lambda-\beta\lambda\right) = \lambda.$$

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

$$v_0 = \frac{-\left(\kappa^2+\lambda(1-\beta)\right)+\sqrt{\left(\kappa^2+\lambda(1-\beta)\right)^2+4\lambda\beta\kappa^2}}{2\beta\kappa^2},$$

and we can calculate g and h given v_0 and the formula above for x and π.

C Proof of proposition for the backward-looking model

The problem is

$$\begin{aligned}
&\min_{x_0}\frac{1}{2}\left(\pi_0^2+\lambda x_0^2\right)+\beta L(x_0,\pi_0),\\
s.t.\quad &:\\
x_0 &\leq \delta x_{-1}+\frac{\rho_0^n+\pi_{-1}}{\sigma},\\
\pi_0 &= \pi_{-1}+\kappa x_0.
\end{aligned}$$

Define $U(x_0,p)=\frac{1}{2}\left((\pi_{-1}+\kappa x_0)^2+\lambda x_0^2\right)+\beta L(x_0,\pi_{-1}+\kappa x_0)$. We wish to prove that the optimal solution to the problem $\min_{x_0} U(x_0,p)$ is increasing in p. Assuming an interior solution, the optimal output gap is characterized by

$$U_1(x_0^*(p),p)=0,$$

and $U_{11}(x_0^*(p), p) > 0$. Differentiating yields:

$$\frac{dx_0^*(p)}{dp} = -\frac{U_{12}(x_0^*(p), p)}{U_{11}(x_0^*(p), p)},$$

so the sign of $\frac{dx_0^*(p)}{dp}$ is the opposite of the sign of $U_{12}(x_0^*(p), p)$. We can calculate

$$\begin{aligned} U_2(x_0, p) &= \frac{1+\beta V}{2}\left(\kappa\delta x_0 + (\pi_{-1} + \kappa x_0)\left(1 + \frac{\kappa}{\sigma}\right) + \kappa\frac{\rho}{\sigma}\right)^2 \\ &\quad + \frac{\lambda}{2}\left(\delta x_0 + \frac{\pi_{-1} + \kappa x_0}{\sigma} + \frac{\rho}{\sigma}\right)^2 - \frac{V}{2}\left(\pi_{-1} + \kappa x_0\right)^2, \end{aligned}$$

and hence

$$\begin{aligned} U_{12}(x_0, p) &= (1+\beta V)\,\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)\left(\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)x_0 + \pi_{-1}\left(1 + \frac{\kappa}{\sigma}\right) + \kappa\frac{\rho}{\sigma}\right) \\ &\quad + \lambda\left(\delta + \frac{\kappa}{\sigma}\right)\left(\left(\delta + \frac{\kappa}{\sigma}\right)x_0 + \frac{\pi_{-1}}{\sigma} + \frac{\rho}{\sigma}\right) - V\kappa\left(\pi_{-1} + \kappa x_0\right), \\ &= (1+\beta V)\,\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)x_0 + (1+\beta V)\,\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)\pi_{-1}\left(1 + \frac{\kappa}{\sigma}\right) \\ &\quad + (1+\beta V)\,\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)\kappa\frac{\rho}{\sigma} + \lambda\left(\delta + \frac{\kappa}{\sigma}\right)\left(\delta + \frac{\kappa}{\sigma}\right)x_0 \\ &\quad + \lambda\left(\delta + \frac{\kappa}{\sigma}\right)\frac{\pi_{-1}}{\sigma} + \lambda\left(\delta + \frac{\kappa}{\sigma}\right)\frac{\rho}{\sigma} - V\kappa\left(\pi_{-1} + \kappa x_0\right), \\ &= \left((1+\beta V)\,\kappa^2\left(1 + \delta + \frac{\kappa}{\sigma}\right)^2 - V\kappa^2 + \lambda\left(\delta + \frac{\kappa}{\sigma}\right)^2\right)x_0 \\ &\quad + \left((1+\beta V)\,\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)\left(1 + \frac{\kappa}{\sigma}\right) - V\kappa + \lambda\left(\delta + \frac{\kappa}{\sigma}\right)\frac{1}{\sigma}\right)\pi_{-1} \\ &\quad + \left(\lambda\left(\delta + \frac{\kappa}{\sigma}\right) + (1+\beta V)\,\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)\kappa\right)\frac{\rho}{\sigma}. \end{aligned}$$

It is immediate that if $\pi_{-1} = 0$ and if $x_0 = 0$, then

$$U_{12}(x_0, p) = \left(\lambda\left(\delta + \frac{\kappa}{\sigma}\right) + (1+\beta V)\,\kappa\left(1 + \delta + \frac{\kappa}{\sigma}\right)\kappa\right)\frac{\rho}{\sigma} < 0.$$

Since $\pi_0 = \pi_{-1} + \kappa x_0$ and $i_0 = \rho_0^n + \pi_{-1} - \sigma(x_0 - \delta x_{-1})$, the other results follow immediately.

D Forward-looking Phillips curve and backward-looking IS curve

In this section, we demonstrate that a result similar to the previous subsection holds even in a model that is partly forward-looking. Specifically, we consider an IS curve with lagged output,[36]

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

but we now use a forward-looking Phillips curve, $\pi_t = \beta E_t\pi_{t+1} + \kappa x_t$; and we assume again that there is only uncertainty about the natural rate of interest; *i.e.* the central bank observes ρ_0^n and u_0; ρ_1^n is uncertain but $u_1 = 0$; and $\rho_t^n = \rho$ and $u_t = 0$ for $t \geq 2$. Because the expected inflation next period affects inflation today, this model will also give rise to the "expectational channel". Hence our results here will mix both channels.

To calculate optimal policy, start at time 2; we know that $u_t = 0$ and $\rho_t^n = \rho$ forever. The state of the economy is also determined by the lagged output gap x_1. However, the only effect of x_1 is to shift the IS curve, which can be offset by appropriate monetary policy, assuming that the ZLB does not bind at time 2, *i.e.* that ρ is large enough. Mathematically, we can set $\pi_t = x_t = 0$ for $t \geq 2$ by setting $i_t = \rho + \sigma\delta x_{t-1}$. The loss function is thus 0 for time $t = 2$ on.

Turn now to optimal policy at time 1. For a given ρ_1^n realization, and given last period's output gap x_0, the problem can be written as:

$$\begin{aligned} V(x_0, \rho_1^n) &= \min_{x_1} \frac{1}{2}\left(\pi_1^2 + \lambda x_1^2\right), \\ s.t. &: \\ \pi_1 &= \kappa x_1, \\ x_1 &\leq \delta x_0 + \frac{\rho_1^n}{\sigma}, \end{aligned}$$

where we have exploited that (i) the future loss is zero and (ii) $E\pi_2 = Ex_2 = 0$. Clearly, the optimum is to set $x_1 = \pi_1 = 0$ if possible. This is feasible if $\delta x_0 + \frac{\rho_1^n}{\sigma} \geq 0$. Hence, the solution:

(i) if $\rho_1^n \geq -\sigma\delta x_0$, set $x_1 = \pi_1 = 0$ with $i_1 = \rho_1^n + \sigma\delta x_0$;

(ii) if $\rho_1^n < -\sigma\delta x_0$, set $x_1 = \delta x_0 + \frac{\rho_1^n}{\sigma}$ and $\pi_1 = \kappa\left(\delta x_0 + \frac{\rho_1^n}{\sigma}\right)$, with $i_1 = 0$.

Overall, the probability of hitting the ZLB is

$$P(x_0) = \int_{-\infty}^{-\sigma\delta x_0} f_\rho(\rho)d\rho,$$

[36] This is a special case of the mixed backward-forward looking equation often used in DSGE models,

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

Assuming no forward-looking term whatsoever simplifies our analysis.

which is again endogenous. The expected loss at time 1 is

$$\begin{aligned} L(x_0) &= \int_{-\infty}^{\infty} V(x_0,\rho) f_\rho(\rho) d\rho \\ &= \frac{\lambda+\kappa^2}{2} \int_{-\infty}^{-\sigma\delta x_0} \left(\frac{\rho}{\sigma} + \delta x_0\right)^2 f_\rho(\rho) d\rho, \end{aligned}$$

which is decreasing and convex: $L' < 0$, $L'' > 0$ and $\lim_{x\to+\infty} L(x) = 0$. Having a higher lagged output gap x_0 is always a positive at time 1 (the loss is decreasing in x_0). If the natural rate shock is high, we simply offset the momentum given by the positive lagged output gap through higher interest rates. But if the natural rate shock is low, a higher lagged output gap helps reduce the severity of the ZLB problem.

We can calculate the expected output gap and inflation at time 1:

$$Ex_1 = \int_{-\infty}^{-\sigma\delta x_0} \left(\delta x_0 + \frac{\rho}{\sigma}\right) f_\rho(\rho) d\rho,$$

and $E\pi_1 = \kappa E x_1$; note that a higher lagged output gap increases these expectations, but that both are always negative.

Given that the loss at time 1, $L(x_0)$, is affected by time 0 choices, the policy maker will set the interest rate (output gap) at time 0 to solve the problem

$$\begin{aligned} &\min_{x_0} \frac{1}{2}\left(\pi_0^2 + \lambda x_0^2\right) + \beta L(x_0) \\ s.t. &: \\ \pi_0 &= \beta E\pi_1 + \kappa x_0 + u_0, \\ x_0 &\leq \delta x_{-1} + \frac{\rho_0^n}{\sigma}. \end{aligned}$$

As in the purely backward-looking problem, the policymaker takes into account that stimulating the economy today (time 0) yields an additional benefit at time 1, as reflected in the function $L(x_0)$; this will naturally lead to a looser (more stimulative) policy. Compared to the previous section, there are two simplifications: first, there is only one endogenous state variable, x_0, rather than two; second, there is no cost tomorrow of too much output gap, which simplifies the loss function. There is however one new complication, that owes to the forward-looking nature of the Phillips curve: the policymaker must now take into account that expected inflation tomorrow depends on the output gap today, since $E\pi_1 = \kappa \int_{-\infty}^{-\sigma\delta x_0} \left(\delta x_0 + \frac{\rho}{\sigma}\right) f_\rho(\rho) d\rho$, which feeds back into inflation today since inflation is forward looking. This means that the effective Phillips curve becomes steeper - a given positive output gap translates into more inflation pressure. The central bank may want to offset this inflation by tightening policy. To find out which effect dominates, take the first-order condition (assuming for now that the ZLB does not bind at time 0), which reads

$$(\kappa + \beta\kappa\delta P(x_0))\left(\beta E\pi_1 + \kappa x_0 + u_0\right) + \lambda x_0 + \beta L'(x_0) = 0,$$

and note that the first term reflects the higher effective Phillips curve slope, which we can define as $\widehat{\kappa} = \kappa\left(1 + \beta\delta P(x_0)\right)$, and the last term $L'(x_0)$ captures the benefits to higher output gap tomorrow. For instance, in the special case where the current cost-push shock just offsets expected inflation, $\beta E\pi_1 + u_0 = 0$, we see that the optimal output gap is determined by

$$\left(\kappa\widehat{\kappa} + \lambda\right) x_0 + \beta L'(x_0) = 0,$$

i.e. we equate the marginal cost of a higher output gap today (more output and inflation today) with the benefit of a higher output gap tomorrow (less severity because of the possibly binding ZLB). In contrast, if the buffer stock channel does not exist (if $\delta = 0$), then it is optimal to set $x_0 = 0$; for instance that is what would happen in the "expectational channel" section if $\beta E\pi_1 + u_0 = 0$.

In general, when $\beta E\pi_1 + u_0 \neq 0$, we can use that $L'(x_0) = \frac{\lambda+\kappa^2}{\sigma^2}\frac{\delta}{\kappa}E\pi_1$ to rewrite the decision rule as:

$$x_0 = -\frac{\widehat{\kappa} + \phi}{\lambda + \kappa\widehat{\kappa}}\beta E\pi_1 - \frac{\widehat{\kappa}}{\lambda + \kappa\widehat{\kappa}}u_0,$$

where ϕ is defined as $\frac{\lambda+\kappa^2}{\sigma^2}\frac{\delta}{\kappa}$. This contrasts with the rule when there is no "momentum" (*i.e.* when $\delta = 0$):

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

Recall that $E\pi_1 < 0$ due to the possibly binding ZLB constraint. It is easy to see that the coefficients in front of $\beta E\pi_1$ and u_0 are larger when $\delta > 0$. This implies that the central bank leans more against expected deflation. And as in the analysis of section 2.1, this loose bias may eventually push it to set $i_0 = 0$, *i.e.* to hit the ZLB *today.*

Of course, this result, while useful, is not a full solution since $E\pi_1$ ultimately depends on x_0. To obtain the full solution, we resort to the same simplification as in the previous section and assume that the distribution of ρ_1^n is such that there is only one negative value of ρ, call it $\underline{\rho}$; it arises with probability p; and the solution will be such that this is the only value of ρ for which the ZLB binds, *i.e.* $\delta x_0 + \frac{\underline{\rho}}{\sigma} < 0$. In this case, we can simplify the expressions of the expected loss function, and the expected output gap:

$$L(x_0) = \frac{\lambda + \kappa^2}{2}p\left(\frac{\underline{\rho}}{\sigma} + \delta x_0\right)^2,$$

as well as $L'(x_0) = \left(\lambda + \kappa^2\right)p\delta\left(\frac{\underline{\rho}}{\sigma} + \delta x_0\right)$ and $Ex_1 = p\left(\delta x_0 + \frac{\underline{\rho}}{\sigma}\right)$. Plugging these in the first-order condition yields, after some simplifications,

$$x_0 = -\frac{\kappa(1+\beta\delta p)}{\kappa^2(1+\beta\delta p)^2 + \lambda + \beta\left(\lambda+\kappa^2\right)p\delta^2}u_0 - \frac{\kappa^2(1+\beta\delta p) + \delta\left(\lambda+\kappa^2\right)}{\kappa^2(1+\beta\delta p)^2 + \lambda + \beta\left(\lambda+\kappa^2\right)p\delta^2}\beta p\frac{\underline{\rho}}{\sigma}.$$

This expression can be usefully compared that obtained in section 2.1,

$$x_0 = -\frac{\kappa}{\kappa^2 + \lambda} u_0 - \frac{\kappa^2}{\kappa^2 + \lambda} \beta p \frac{\underline{\rho}}{\sigma}.$$

First, notice that when $\delta = 0$, the two formulas coincide. Second, one can show that the optimal policy x_0 is increasing in p, for p small enough. And third, one can show that, x_0 is more increasing in p when δ is large (at least if $u_0 = 0$, p and δ are small). Our analysis is summarized in:

Proposition 6 *When the IS curve is backward-looking, optimal policy becomes looser when the risk of hitting the ZLB becomes larger, or when the severity of the ZLB becomes bigger. This effect is more pronounced when there is more persistence in the IS curve (*i.e. *when δ is larger). Technically, x_0 and π_0 are increasing in p and decreasing in $\underline{\rho}$, while i_0 is decreasing in p and increasing in $\underline{\rho}$; and these rates of variation are larger when δ is larger.*

The proof obtains from studying the function

$$G(p, \delta) = \frac{\kappa^2(1 + \beta\delta p) + \delta\left(\lambda + \kappa^2\right)}{\kappa^2(1 + \beta\delta p)^2 + \lambda + \beta\left(\lambda + \kappa^2\right) p\delta^2} \beta p \frac{\underline{\rho}}{\sigma}$$

and showing that G is increasing in p for small p, and G_1 is increasing in δ for small p and δ. (Details to be added.)

E Backward-looking Phillips curve with a forward-looking IS curve

In this section, we study the case where the Phillips curve is backward-looking, but the IS curve is forward-looking. The basic intuition is similar to the case with a forward-looking Phillips curve, and backward-looking IS curve, in that a buffer stock channel arises naturally. However, there are a couple of important differences. First, as in the case of a purely backward-looking model, there is a loss of having too high inflation if the ZLB constraint does not bind tomorrow (whereas, with a backward-looking IS curve, the extra output gap can always be eliminated by an appropriate choice of the interest rate if the ZLB does not bind). Second, a complication emerges because inflation today now affects inflation tomorrow, and this feeds backward to affect output today (both directly and by affecting output tomorrow). This makes it more difficult to analyze this case.

The Phillips curve is

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

and the IS curve is

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

We focus here again on the case where the uncertainty is about ρ_1^n, and it is known in advance that $u_t = 0$ for $t \geq 1$.

First, note that substituting the Phillips curve in the IS curve yields a simplified equation,

$$x_t = E_t x_{t+1}\left(1+\frac{\kappa}{\sigma}\right) - \frac{1}{\sigma}\left(i_t - \rho_t^n - \pi_t\right).$$

To find optimal policy, we again start at time 2. The economy does not experience any more shock, so $\rho_t^n = \rho$, but it starts with an initial "inflation momentum" π_1 (which may be positive or negative). As in the backward-looking model, the optimal policy under discretion involves reducing this inflation back to target, which requires an output gap in the opposite direction. Given the quadratic loss, is optimal to smooth this adjustment over time. It can be shown that optimal policy under discretion yields a value function $\frac{V}{2}\pi_1^2$, and optimal policies are linear in the lagged inflation, i.e. they can be represented as $\pi_t = g\pi_{t-1}$ and $x_t = h\pi_{t-1}$ where g and h are constants. (Details to be added; cite Jeff Campbell's paper on existence and uniqueness of the equilibrium.) We can prove that $g \in [0,1]$ and $h < 0$.

We can now turn to the optimal policy at time 1. The policymaker problem is, given a starting value π_0 and a shock realization ρ_1^n, to decide x_1 subject to the ZLB constraint, taking into account that the inflation today will be costly tomorrow. A first step in this solution is to rewrite the ZLB constraint:

$$x_1 \leq E x_2\left(1+\frac{\kappa}{\sigma}\right) + \frac{1}{\sigma}\left(\rho_1^n + \pi_1\right),$$

as

$$x_1 \leq h\pi_1\left(1+\frac{\kappa}{\sigma}\right) + \frac{1}{\sigma}\left(\rho_1^n + \pi_1\right),$$

and given that $\pi_1 = \pi_0 + \kappa x_1$, we can further simplify this to

$$x_1 \leq \frac{1}{1-\kappa\left(h\left(1+\frac{\kappa}{\sigma}\right)+\frac{1}{\sigma}\right)}\left(\left(h\left(1+\frac{\kappa}{\sigma}\right)+\frac{1}{\sigma}\right)\pi_0 + \frac{1}{\sigma}\rho_1^n\right),$$

which we rewrite as

$$x_1 \leq \zeta_0\pi_0 + \zeta_1\rho_1^n.$$

We will assume that $\zeta_0 > 0$ and $\zeta_1 > 0$. (We should be able to prove that this is true at least under some parameter restrictions.) That is, a higher inflation or natural rate today reduces the likelihood that the ZLB will bind.[37]

Taken this new ZLB constraint into account, the problem at time 1 is

$$\begin{aligned} V(\pi_0, \rho_1^n) &= \min_{x_1} \frac{1}{2}\left(\pi_1^2 + \lambda x_1^2\right) + \beta\frac{V}{2}\pi_1^2 \\ s.t. &: \end{aligned}$$

[37] The reason this is not a priori obvious is that higher inflation today yields lower output tomorrow according to the policy function h, which tightens the ZLB constraint.

$$\begin{aligned}\pi_1 &= \pi_0 + \kappa x_1, \\ x_1 &\leq \zeta_0 \pi_0 + \zeta_1 \rho_1^n.\end{aligned}$$

There are two cases. Either the ZLB does not bind, in which case optimal policy is determined by

$$x_1 = -\frac{(1+\beta V)\,\kappa}{(1+\beta V)\kappa^2 + \lambda}\pi_0,$$

and

$$\pi_1 = \frac{\lambda}{(1+\beta V)\kappa^2 + \lambda}\pi_0,$$

with loss

$$V(\pi_0, \rho_1^n) = \frac{1}{2}V\pi_0^2,$$

since this is the same problem as at time 2. Or the ZLB binds, and then

$$\begin{aligned}x_1 &= \zeta_0 \pi_0 + \zeta_1 \rho_1^n, \\ \pi_1 &= \pi_0 + \kappa x_1 = \pi_0(1+\kappa\zeta_0) + \kappa\zeta_1\rho_1^n,\end{aligned}$$

and the loss in this case is

$$\begin{aligned}V(\pi_0, \rho_1^n) &= \frac{1}{2}\left((1+\beta V)\pi_1^2 + \lambda x_1^2\right) \\ &= \frac{1}{2}\left((1+\beta V)\left(\pi_0(1+\kappa\zeta_0) + \kappa\zeta_1\rho_1^n\right)^2 + \lambda\left(\zeta_0\pi_0 + \zeta_1\rho_1^n\right)^2\right).\end{aligned}$$

The threshold value for a binding ZLB is

$$\rho^*(\pi_0) = -\left(\frac{(1+\beta V)\,\kappa}{(1+\beta V)\kappa^2 + \lambda} + \zeta_0\right)\frac{\pi_0}{\zeta_1},$$

which is decreasing in π_0 under our assumptions that $\zeta_0 > 0$ and $\zeta_1 > 0$. Hence, the probability of hitting the ZLB is

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

and is higher when π_0 is lower. The total expected loss at time 1 is

$$\begin{aligned}L(\pi_0) &= \int_{-\infty}^{\infty} V(\pi_0, \rho) f_\rho(\rho)d\rho, \\ &= \int_{-\infty}^{\rho^*(\pi_0)} \frac{1}{2}\left((1+\beta V)\left(\pi_0(1+\kappa\zeta_0) + \kappa\zeta_1\rho\right)^2 + \lambda\left(\zeta_0\pi_0 + \zeta_1\rho\right)^2\right) f_\rho(\rho)d\rho \\ &\quad + \frac{1}{2}V\pi_0^2 \int_{\rho^*(\pi_0)}^{\infty} f_\rho(\rho)d\rho.\end{aligned}$$

The structure of the loss is very similar to the structure in the purely backward-looking model studied in the main body of the paper. Here too, a marginal change in ρ^* does not affect the total expected loss. This suggests it is useful to focus, as in our analysis of the backward-looking model, on the simple scenario where ρ can take only one negative value, which is the only one where the ZLB binds. Call $\underline{\rho}$ the low value of ρ and p its probability. Then

$$L(\pi_0) = \frac{V}{2}\pi_0^2 + p\left(\frac{1}{2}\left((1+\beta V)\left(\pi_0(1+\kappa\zeta_0)+\kappa\zeta_1\underline{\rho}\right)^2 + \lambda\left(\zeta_0\pi_0+\zeta_1\underline{\rho}\right)^2\right) - \frac{V}{2}\pi_0^2\right),$$

and the time 0 problem is, given π_{-1}, u_0 and ρ_0^n, to solve:

$$\begin{aligned} & \min_{x_0} \frac{1}{2}\left(\pi_0^2 + \lambda x_0^2\right) + \beta L(\pi_0), \\ s.t. \quad & : \\ \pi_0 \quad & = \quad \pi_{-1} + \kappa x_0 + u_0, \\ x_0 \quad & \leq \quad \zeta_0\pi_{-1} + \zeta_1\rho_0^n. \end{aligned}$$

The main result that we conjecture is the following:

Proposition 7 *Suppose the Phillips curve is backward-looking and the IS curve forward-looking. Suppose that $\pi_{-1} + u_0 = 0$. Then x_0 and π_0 are increasing in p, and i_0 is decreasing in p (at least for p small enough).*

The proof works similarly to that of the backward-looking model. Define

$$U(x_0, p) = \frac{1}{2}\left((\pi_{-1} + \kappa x_0 + u_0)^2 + \lambda x_0^2\right) + \beta L(\pi_{-1} + \kappa x_0 + u_0),$$

and note that the optimal solution satisfies

$$\begin{aligned} U_1(x_0^*(p), p) \quad & = \quad 0, \\ U_{11}(x_0^*(p), p) \quad & > \quad 0. \end{aligned}$$

We have $\frac{dx_0^*}{dp} = -\frac{U_{12}(x_0^*(p),p)}{U_{11}(x_0^*(p),p)}$ so the sign of the comparative statics depends on

$$U_{12}(x_0, p) = \left((1+\beta V)\,\kappa(1+\kappa\zeta_0)\left(\kappa x_0(1+\kappa\zeta_0)+\kappa\zeta_1\underline{\rho}\right) + \zeta_0\kappa\lambda\left(\zeta_0\kappa x_0+\zeta_1\underline{\rho}\right)\right) - V\kappa^2 x_0,$$

and note that

$$U_{12}(0, p) = \left((1+\beta V)\,\kappa(1+\kappa\zeta_0)\kappa\zeta_1 + \zeta_0\kappa\lambda\zeta_1\right)\underline{\rho}$$

which is positive, again under our assumption that $\zeta_0 > 0$ and $\zeta_1 > 0$.