Appendices for
"Risk Management for Monetary Policy
Near the Zero Lower Bound"*

Charles Evans   Jonas Fisher   François Gourio   Spencer Krane

May 21, 2015

A Theoretical details

This appendix presents the several theoretical results that are referred to in the main text.

A.1 Optimal policy in the forward-looking model with uncertainty about cost-push shocks

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

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

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

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

*Any views expressed herein are those of the authors and do not necessarily represent those of the Federal Open Market Committee or the Federal Reserve System.

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

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

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

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

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 1 *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.

A.2 Calculation of W in the backward-looking model

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

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

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

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

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

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

The first order condition yields

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

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

leading to

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

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

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

which can be simplified to a simple quadratic equation:

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

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

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

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

A.3 Proof of Proposition 2

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

Lemma 1 *Consider the problem*

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

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

$$\begin{aligned} J(x_0, \varepsilon) &= \max_{x_1} F(x_1, x_0, \varepsilon), \\ s.t. &: x_1 \le f(x_0) + \varepsilon, \end{aligned}$$

where F is quadratic (with $F_{11} < 0$) and f is linear. Suppose that higher θ indexes more risky distribution of ε in the sense of second-order stochastic dominance. Suppose that the scalar $F_{13} + F_{11} < 0$ and that the scalar $f'(F_{11} + F_{13}) + F_{21}\left(1 + \frac{F_{13}}{F_{11}}\right) < 0$. Then, x_0 is increasing in θ.

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

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

It is straightforward from the implicit function theorem that

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

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

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

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

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

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

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

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

and using the envelope theorem we calculate

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

Since F is quadratic and f is linear, (and hence x_1 is linear), the two expressions for J_1 are both linear in ε. To determine the convexity of this function simply requires comparing the slopes.[1]

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

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

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

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

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

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

■

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

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

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

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

and the ZLB constraint can be rewritten as

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

or

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

where

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

It is thus possible to rewrite the Bellman equation as

$$\begin{aligned}
W_t(\pi_{t-1}, \overline{\pi}_t, \rho, u) &= \min_{\pi_t} \frac{1}{2}\left(\pi_t^2 + \frac{\lambda}{\kappa^2}\left(\pi_t - \xi\pi_{t-1} - u\right)^2\right) + \beta E_{\rho'|\rho} W_{t+1}(\pi_t, \overline{\pi}_{t+1}, \rho', u'), \\
s.t. &: \\
\pi_t &\leq \overline{\pi}_t, \\
\overline{\pi}_{t+1} &= \left(\xi + \delta + \frac{\kappa}{\sigma}\right)\pi_t + \frac{\kappa}{\sigma}\rho' - \xi\delta\pi_{t-1} - \delta u + u'.
\end{aligned}$$

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

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

This value function will of course be quadratic:

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

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

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

Finally, the time 0 problem is

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

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

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

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

and

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

The problem is then

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

where

$$\begin{aligned} J(\pi_0, \varepsilon) &= \max_{\pi_1} F(\pi_1, \pi_0, \varepsilon) \\ s.t. &: \pi_1 \leq f(\pi_0) + \varepsilon, \end{aligned}$$

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

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

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

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

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

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

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

Second,

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

B Forward-looking model solution methods

We present here the numerical methods used to solve the forward-looking model. We make the following assumptions regarding exogenous variables. First, there is a date T such that, for $t \geq T$, the cost-push shock is zero and the natural rate is constant, $u_t = 0$ and $\rho_t^n = \overline{\rho}$. Second, for $t < T$, the cost-push shock u_t follows a Markov chain with transition probability $P_u(u'|u)$. The natural rate ρ_t^n is the sum of a deterministic component and a Markov chain:

$\rho_t^n = f_t + \varepsilon_t$, where ε_t has transition $P_\varepsilon(\varepsilon'|\varepsilon)$, and f_t is increasing and satisfies $f_T = \overline{\rho}$. We will write $\rho_t^n(\varepsilon) = f_t + \varepsilon$. The stochastic processes ε_t and u_t are independent. In practice we use simply $f_t = \rho_0^n + \frac{t}{T_0}(\overline{\rho} - \rho_0^n)$ for $0 \le t \le T_0$ and $f_t = \overline{\rho}$ for $T_0 \le t < T$. We will choose the Markov chains for ε and for u to each approximate an AR(1) process using the Rouwenhorst method.

The model we study is

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

Our theoretical analysis assumed for simplicity (and as is common in the literature) a zero inflation steady-state. To provide more useful numerical illustrations, we consider the case of a positive inflation target. We assume that the equations above apply if π_t is inflation deviation from target and i_t is the nominal rate minus the inflation target. The ZLB is then modified as $i_t \ge \overline{Z} \overset{\text{def}}{=} -\pi^*$.[2]

B.1 Calculation of optimal policy under discretion

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

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

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

Denote

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

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

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

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

and

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

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

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

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

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

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

In that case, the solution is:

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

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

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

and

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

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

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

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

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

B.2 Calculation of equilibrium under a Taylor rule

Suppose that the central bank follows the policy:

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

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

- Set i_t equal to the current natural real rate of interest, $i_t = \max\left(\overline{Z}, \rho_t^n(\varepsilon_t)\right)$; this corresponds to $g_t(\varepsilon_t) = \rho_t^n(\varepsilon_t)$, and $\phi = \gamma = 0$;
- A Taylor rule with a constant intercept, $i_t = \max\left(\overline{Z}, \widehat{\rho} + \phi\pi_t + \gamma x_t\right)$.

The system of equations to solve is

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

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

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

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

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

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

and note that

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

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

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

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

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

and we check that with this solution, $g_t(\varepsilon_t) + \phi\pi_t^W(\varepsilon, u) + \gamma x_t^W(\varepsilon, u) < \overline{Z}$.[3] Given the value of $\pi_t^W(\varepsilon, u)$ and $x_t^W(\varepsilon, u)$ for all ε, u, we can update $a_{t-1}^W(\varepsilon, u)$ and $b_{t-1}^W(\varepsilon, u)$ and hence proceed backwards until time 0. We can furthermore calculate the loss function in the same way as for optimal policy.

B.3 Calculation of equilibrium under "Naive" Policy

Since the Fed does not recognize the possibility of future shocks, in any given period t, starting with shocks ρ_t^n and u_t, the Fed assumes that in the future, the natural rate and the cost push shock will simply revert deterministically to their respective trends. This implied path for $\left\{ \rho_{t+k}^n, u_{t+k} \right\}_{k=1}^{T-t}$ is calculated using the true persistence of the transitory shocks (ε, u) and the true deterministic trend of the natural real rate. The Fed then sets the nominal interest rate i_t to minimize the loss $\sum_{k=0}^{T-t} \beta^k \left(\pi_{t+k}^2 + \lambda x_{t+k}^2 \right)$, subject to

$$\begin{aligned} x_{t+k} &= x_{t+k+1} - \frac{1}{\sigma} \left(i_{t+k} - \rho_{t+k}^n - \pi_{t+k+1} \right), \\ \pi_{t+k} &= \beta \pi_{t+k+1} + \kappa x_{t+k} + u_{t+k}, \end{aligned}$$

and assuming discretion. Hence, at each point in time, the Fed solves the backward induction problem from time T on, for a given path $\left\{ \rho_{t+k}^n, u_{t+k} \right\}_{k=1}^{T-t}$, and deduces the optimal interest rate today $i_t = i_t(\varepsilon, u)$.

Once this rule has been calculated, we can then use it to solve for the behavior of private agents. This implies that we assume the agents understand that the Fed will behave naively in the future.

If the Fed follow this policy, it is then surprised each period in two ways: first and most obviously, it is surprised by the realization of new shocks each period. Second, after setting the interest rate, the Fed is surprised by the realized value of x_t and π_t. This is because the

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

Fed has used its projected values for x_{t+1} and π_{t+1} in the model, while the true values are agent's expectations $E_t x_{t+1}$ and $E_t \pi_{t+1}$. Of course, if there was no uncertainty in reality, or if the zero lower bound was never binding, the two would coincide.

C Backward-looking model solution methods

The backward-looking model is

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

and

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

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

C.1 Calculation of optimal policy under discretion

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

$$\begin{aligned}
V_t(x_{-1}, \pi_{-1}, \varepsilon, u) &= \min_{i \geq \overline{Z}} \frac{1}{2}\left(\pi^2 + \lambda x^2\right) + \beta E_{\varepsilon', u' | \varepsilon, u} V_{t+1}(x, \pi, \varepsilon', u'), \\
s.t. &: \\
\pi &= \xi \pi_{-1} + \kappa x + u, \\
x &= \delta x_{-1} - \frac{1}{\sigma}\left(i - \rho_t^n(\varepsilon) - \pi_{-1}\right).
\end{aligned}$$

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

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

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

C.2 Equilibrium under a Taylor rule

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

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

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

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

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

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

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

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

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

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

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

C.3 Calculation of equilibrium under "Naive" Policy

We also compute a counterfactual where the Fed behaves "naively" in the sense that it optimizes but does not recognize the possibility of future shocks. This section details our computation of this counterfactual in the backward-looking model.

Similar to the forward-looking model, the naive Fed does not recognize the possibility of future shocks, and in each state $t, \rho_t^n, u_t, x_{t-1}, \pi_{t-1}$, it assumes a path for $\left\{\rho_{t+k}^n, u_{t+k}\right\}_{k=0}^{T-t}$.

The Fed then solves the optimal stabilization problem without uncertainty, which can be represented with the Bellman equation:

$$\begin{aligned}
\widetilde{V}_t(x_{-1}, \pi_{-1}) &= \min_{i \geq \overline{Z}} \frac{1}{2}\left(\pi^2 + \lambda x^2\right) + \beta \widetilde{V}_{t+1}(x, \pi), \\
s.t. &: \\
\pi &= \xi \pi_{-1} + \kappa x + u_t, \\
x &= \delta x_{-1} - \frac{1}{\sigma}\left(i - \rho_t^n - \pi_{-1}\right),
\end{aligned}$$

and the key difference is that the path for $\left\{\rho_{t+k}^n, u_{t+k}\right\}_{k=0}^{T-t}$ is now fixed. Solving this problem (again backwards from time T) yields a policy function at time t, call it $i_t(x_{-1}, \pi_{-1}, \rho_t^n, u_t)$. Once we have solved for this policy at each point in time and for each possible value of the states, we can again simulate the model given this interest rate rule.

There are two important differences with the forward-looking naive policy. First, agents' expectations are not relevant and hence whether they assume the Fed is naive or not does not feed back on their decisions except through the interest rate. Second, the Fed is not surprised by the achieved levels of output gap and inflation given its interest rate.

C.4 Deflationary traps

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

D Effects of assuming alternative parameter values in the simulations

This appendix summarizes various perturbations of our simulations based on changing one parameter at a time. These simulations are designed to illustrate the effects of key parameters on our findings.

D.1 Forward-looking model

For the forward-looking model we focus on changes in the initial natural real interest rate ρ_0 and the unconditional volatility in the random component of the natural rate, σ_ε. We considered five scenarios summarized in Table 1. The results for the forward-looking model under optimal discretion, the naive policy and the Taylor rule are reported in Tables 2, 3 and 4. Column 1 is the baseline parameterization considered in the main text.

Table 1: Alternative parameters in the forward-looking model

	Perturbations				
Parameter	1	2	3	4	5
ρ_0	-0.50	-0.50	-0.50	-1.00	0.00
σ_ε	2.50	2.00	3.00	2.50	2.50

Table 2: Forward-looking perturbations: Optimal Discretion

Statistic	1	2	3	4	5
Expected loss	0.02	0.02	0.03	0.03	0.02
Mean time at liftoff	4.13	1.00	5.74	5.27	1.00
Median time at liftoff	3	1	4	4	1
Median π at liftoff	-0.08	0.04	-0.88	-0.03	0.06
Median x at liftoff	2.19	1.90	1.38	1.57	1.84
75^{th} percentile $\max(\pi)$	2.77	2.77	2.75	2.38	2.66
25^{th} percentile $\min(x)$	-0.31	-0.31	-1.44	-0.15	-1.22
Median standard deviation Δi	1.85	1.77	1.94	1.77	1.91

Table 3: Forward-looking perturbations: Naive

Statistic	1	2	3	4	5
Expected loss	0.06	0.03	0.11	0.07	0.05
Mean time at liftoff	1.00	1.00	1.00	1.00	1.00
Median time at liftoff	1	1	1	1	1
Median π at liftoff	-1.44	-0.84	-2.20	-1.68	-1.23
Median x at liftoff	0.88	1.36	0.26	0.71	1.02
75^{th} percentile max(π)	2.77	2.77	2.72	2.38	2.34
25^{th} percentile min(x)	-1.44	-0.97	-2.54	-1.68	-1.48
Median standard deviation Δi	1.88	1.79	1.99	1.83	1.92

Table 4: Forward-looking perturbations: Taylor rule

Statistic	1	2	3	4	5
Expected loss	0.16	0.13	0.19	0.20	0.12
Mean time at liftoff	1.00	1.00	1.00	3.54	1.00
Median time at liftoff	1	1	1	3	1
Median π at liftoff	-1.62	-1.57	-1.69	-1.73	-1.23
Median x at liftoff	0.35	0.38	0.30	0.55	0.73
75^{th} percentile max(π)	3.10	2.82	2.68	3.04	2.09
25^{th} percentile min(x)	-1.78	-2.08	-3.27	-2.10	-3.15
Median standard deviation Δi	0.97	0.90	1.05	0.94	0.99

D.2 Backward-looking model

For the backward-looking model we focus on changes in the output gap persistence, δ in the IS curve, inflation persistence in the Phillips curve, ξ, initial natural real interest rate ρ_0 and the unconditional volatility in the random component of the natural rate, σ_ε, the initial output gap x_0 and the initial inflation rate π_0. We considered ten scenarios summarized in Table 5. The results for the backward-looking model under optimal discretion, the naive policy and the Taylor rule are reported in Tables 6, 7 and 8. Column 1 is the baseline parameterization considered in the main text.

Table 5: Alternative parameterizations in the backward-looking model

	Perturbations									
Parameter	1	2	3	4	5	6	7	8	9	10
δ	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.65	0.85	0.75
ξ	0.95	0.95	0.95	0.95	0.95	0.93	0.97	0.95	0.95	0.95
ρ_0	-0.50	-0.50	-0.50	-4.00	0.00	-0.50	-0.50	-0.50	-0.50	-0.50
$\sigma(\varepsilon)$	2.50	2.00	3.00	2.50	2.50	2.50	2.50	2.50	2.50	2.50
x_0	-1.50	-1.50	-1.50	-1.50	-1.50	-1.50	-1.50	-1.50	-1.50	0.00
π_0	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70

Table 6: Backward-looking perturbations: Optimal Discretion

Statistic	1	2	3	4	5	6	7	8	9	10
Expected loss	0.27	0.25	0.30	0.69	0.24	0.20	0.59	0.23	0.63	0.22
Mean time at liftoff	12.5	12.2	12.8	20.1	11.3	10.9	14.1	11.6	13.6	9.8
Median time at liftoff	10	10	11	22	9	9	12	9	11	7
Median π at liftoff	0.32	0.28	0.37	0.00	0.36	0.22	0.45	0.28	0.43	0.40
Median x at liftoff	2.00	2.00	2.00	2.00	1.99	1.92	2.00	2.00	2.00	2.00
75^{th} percentile $\max(\pi)$	3.02	3.02	3.03	2.63	3.05	3.01	3.05	3.02	3.03	3.12
25^{th} percentile $\min(x)$	-1.66	-1.58	-1.77	-2.18	-1.66	-1.54	-1.79	-1.65	-1.73	-1.72
Median standard deviation Δi	2.85	2.82	2.96	1.56	3.07	2.84	2.90	2.92	2.77	3.20

Table 7: Perturbations: Naive

Statistic	1	2	3	4	5	6	7	8	9	10
Expected loss	0.28	0.26	0.31	0.69	0.25	0.20	0.60	0.24	0.64	0.22
Mean time at liftoff	10.2	10.1	10.3	18.8	9.14	9.36	11.2	9.36	11.5	6.53
Median time at liftoff	7	7	7	21	6	7	8	6	9	3
Median π at liftoff	0.00	-0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.12
Median x at liftoff	1.83	1.83	1.81	2.00	1.78	1.80	1.86	1.81	1.82	1.79
75^{th} percentile $\max(\pi)$	2.83	2.84	2.83	2.53	2.86	2.87	2.79	2.84	2.83	2.90
25^{th} percentile $\min(x)$	-1.71	-1.64	-1.81	-2.18	-1.72	-1.61	-1.82	-1.71	-1.77	-1.77
Median standard deviation Δi	3.00	2.98	3.11	1.64	3.24	3.03	3.00	3.14	2.86	3.33

Table 8: Backward-looking perturbations: Taylor

Statistic	1	2	3	4	5	6	7	8	9	10
Expected loss	0.60	0.55	0.66	1.22	0.55	0.39	1.53	0.49	1.73	0.53
Mean time at liftoff	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Median time at liftoff	1	1	1	1	1	1	1	1	1	1
Median π at liftoff	-1.27	-1.27	-1.27	-1.67	-1.21	-1.27	-1.26	-1.13	-1.40	-0.23
Median x at liftoff	1.21	1.21	1.21	1.17	1.21	1.22	1.19	1.22	1.19	1.31
75^{th} percentile $\max(\pi)$	2.81	2.78	2.89	2.12	2.93	2.89	2.74	2.94	2.67	2.98
25^{th} percentile $\min(x)$	-1.55	-1.36	-1.75	-2.44	-1.44	-1.54	-1.55	-1.27	-2.00	-1.48
Median standard deviation Δi	0.54	0.54	0.54	0.50	0.54	0.55	0.53	0.54	0.53	0.55

E Data

This appendix describes the sources of the data we use in our empirical analysis.

E.1 Narrative Analysis and Human-coded FOMC-based variables

The narrative analysis is based on Federal Reserve documents, including monetary policy reports, minutes of FOMC meetings, and Records of Policy Actions. These documents are all available through links at the web site:

`http://www.federalreserve.gov/monetarypolicy/default.htm`.

The documentation describing how the human-coded variables were constructed from the Records of Policy Actions and the portions of the FOMC minutes that describe the rationale for the Committee's policy action is in a separate document:

`EFGK-2015-Human-Coding-Appendix.pdf`.

E.2 Machine-coded FOMC-based variables

We searched the policy portion of the FOMC minutes for occurrences of a set of words (risk terms) that appear in the same sentence as a second set of words (conditional terms). The risk terms we use are to capture discussions associated with uncertainty or insurance. The conditional terms relate to economic activity or inflation. We count the number of sentences that include risk-conditional term pairs.

Using plain text files provided by Michael McMahon, we count sentences within the portion of the FOMC minutes that address the policy decision after 1993. From the fifth meeting of 1987 (the first meeting chaired by Alan Greenspan) up to the last meeting of 1992 we use the "Report of Policy Action." The latter is obtained from the Board of Governors web site. Within the text files for the minutes are tags for different sections of the meetings. The policy discussion follows the tag >>FOMC2>>. There were some bugs in the text files we received. These were corrected as follows:

- Several tags were incorrectly marked as >> FOMC2<< or >>FOMC2>>>. These were changed to >>FOMC2>>.
- The transcript for the September 2003 meeting was combined into one file to conform with how all other two-day meeting minutes were formatted.
- The 3/5/1997 and 5/20/1997 files had their names interchanged since their contents stated that they were the minutes for the other date.
- The >>FOMC2>> tag was moved in the 3/31/1998 and 11/13/1999 minutes files to capture the whole policy discussion.

For each unit of text to be searched, the text must first be broken into sentences. We accomplished this using a sentence "tokenizer" from the Natural Language Toolkit.[4] A tokenizer is an algorithm that can distinguish between periods marking the end of a sentence from those that mark abbreviations. While this approach is based on machine learning that could introduce error into the output, we have verified that it is accurate in our case. Once the text was broken into sentences, the text searches proceeded by searching for risk and conditional terms within each sentence. The search terms we use are listed below. Note that the terms are not searched on a case-sensitive basis.

Risk Terms

- Insurance terms: insurance (when not preceded by unemployment, deposit, health, medical, casualty, Federal, life, auto, fire, flood, drought, company, companies, industry, or fund) risk-management, risk management, ensure, and assurance.
- Uncertainty: uncertainty, uncertainties, uncertain, and question.

Conditional terms

- Inflation: inflation, prices, deflation, disinflation, labor cost(s), unit cost(s)
- Activity: activity, growth, slack, resource, labor (when not followed by cost), employment

E.3 Other data

We now describe the remaining data used in our econometric analysis. Some of the data was obtained from the Haver Analytics database. The mnemonic's are given below in these cases.

E.3.1 Fed Funds Rate

We use two different Federal Funds Rate variables. The first is a thirty-day forward average of the target rate following each FOMC meeting (Haver mnemonic: FFEDTAR@DAILY). The second variable uses two different methodologies depending on the date. For meetings prior to 1990 we use the target rate as given in Thornton (2005). Target values for 1990 and later are given by the New York Fed's "Historical Changes of the Target Federal Funds and Discount Rates" available at:

`http://www.newyorkfed.org/markets/statistics/dlyrates/fedrate.html`.

[4] See `www.nltk.org`. For more information on the sentence tokenizer, see `http://www.nltk.org/api/nltk.tokenize.html` or `http://www.nltk.org/book/ch03.html`.

E.3.2 Credit spread

The spread variable (SPD) is calculated as the difference between Moody's Seasoned Baa Corporate Bond Yield percentage points per annum (Haver mnemonic: FBAA@USECON) and the 10-year Treasury Note Yield at Constant Maturity percentage points per annum (Haver mnemonic: FCM10@USECON)

E.3.3 Market Volatility

The Chicago Board Options Exchange Volatility Index (VXO) is based on the prices of eight S&P 100 index put and call options. This measure is also known as the "Original Vix" (Haver mnemonic: SPVXO@WEEKLY). We use this rather than the newer version because it extends back to the beginning of our sample. The two series are highly correlated.

E.3.4 Macroeconomic Uncertainty

Quarterly averages of monthly 12-month ahead macroeconomic uncertainty are calculated using the data and methodology of Jurado, Ludvigson, and Ng (2015). We also create three new uncertainty measures using subsets of their publicly available data: "activity," "inflation," and "other," where "other" is calculated using the residual variables that were unused after calculating activity and inflation uncertainty. All measures are normalized when used in policy rule calculations

- lunc (JLN): macroeconomic uncertainty (all 132 variables)
- luncact: activity uncertainty (70 variables relevant to activity)
- luncinf: inflation uncertainty (24 variables relevant to inflation)
- luncoth: other uncertainty (38 remaining variables)

E.3.5 Federal Open Market Committee "Greenbook" data

Output gap and core CPI data were downloaded from the Philadelphia Federal Reserve's website. This data is from the Board staff's forecast prepared for each FOMC meeting. For quarterly data, values corresponding to the first, third, fifth, and seventh FOMC meetings of each year make up the first, second, third, and fourth quarter values, respectively. The output gap data is obtained from:

`https://www.philadelphiafed.org/research-and-data/real-time-center/greenbook-data/gap-and-financial-data-set.cfm`.

The core CPI data is obtained from:

`https://www.philadelphiafed.org/research-and-data/real-time-center/greenbook-data/philadelphia-data-set.cfm`.

The four-quarter ahead estimate was constructed for each variable by calculating the simple average of the forecasted values for the current quarter and three subsequent quarters, *e.g.*:

$$outputgap\_4q = \frac{outputgap\_T0 + outputgap\_T1 + outputgap\_T2 + outputgap\_T3}{4}$$

where $outputgap\_T0$ is the forecasted output gap for the current quarter, $outputgap\_T1$ is the forecasted output gap for the next quarter, *etc.* A similar expression holds for the $corecpi\_4q$. The variables we use in our empirical analysis are then:

- fcGap: outputgap_4q
- fcInf: corecpi_4q

We construct the interest rate implied by the Taylor (1993) rule (discussed in Sections 2.3.2 and 2.3.3) using the current quarter estimates of core CPI (corecpi_T0) and the output gap (outputgap_T0) from the Greenbook.

The forecast revision variables were calculated as follows. First we calculate the lagged forecast:

$$outputgap\_4q\_TM1 = \frac{L.outputgap\_T1 + L.outputgap\_T2 + L.outputgap\_T3 + L.outputgap\_T4}{4}$$

In this expression "*L.*" denotes "lagged value of." Some consecutive FOMC meetings occur in the same quarter. In these cases

$$outputgap\_4q\_TM1 = \frac{L.outputgap\_T0 + L.outputgap\_T1 + L.outputgap\_T2 + L.outputgap\_T3}{4}$$

The the output gap revision variable (a similar expression holds for the core cpi revision variable) for each meeting is

$$gap\_revision = outputgap\_4q - outputgap\_4q\_TM1$$

The revision variables are:

- frGap: outputgap_revision
- frInf: corecpi_revision

E.3.6 Survey of Professional Forecasters

We use individual point forecasts and bin-based probability forecasts from the Philadelphia Fed's Survey of Professional Forecasters (SPF). These forecasts are contained in Excel files that are found along with documentation at the following web address:

https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/historical-data/individual-forecasts.cfm.

We use the Excel files containing forecasts from the 1980s, 1990s, and 2000s. Relevant worksheet names are given below, along with each variable's description from the SPF documentation.

- RGDP: Quarterly level of real GDP
- PGDP: Quarterly level of the GDP price index
- PRGDP: Probability that annual-average over annual-average percent change in GDP falls in a particular range
- PRPGDP: Probability that annual-average over annual-average percent change in the GDP price index falls in a particular range

The point forecasts and binned forecasts are handled differently due to how the data are collected. Real GDP and GDP deflator point forecasts are given as levels in the SPF. One-year constant horizon growth rates for real GDP, RGDP_ch, are calculated using the following formula:

$$RGDP\_ch = ((RGDP5 - RGDP1)/RGDP1) * 100,$$

where RGDP1 is the historical value for the quarter prior to the survey and RGDP5 is the 3-quarter ahead forecast. See page 13 of the SPF documentation for a description of forecast horizons and page 14 for a table with examples.

For the bin-based probability forecasts we need to construct our own measures of constant horizon forecasts. For each quarterly survey the probability distributions are collected corresponding to forecasts for the current year and the following year. We convert these into probability distributions for forecasts that are one year ahead from the quarter the survey was conducted using the procedure in D'Amico and Orphanides (2014). In particular, for each bin i,

$$probability_i^{c.h.} = \omega_t * current\_year\_probability_i + (1 - \omega_t) * next\_year\_probability_i$$

where $\omega_t = 1.125 - 0.25 * t$ and t is the quarter in which the survey is conducted.

Summary statistics of individual forecaster binned probability distributions were calculated using the assumption that each bin's midpoint is the "point value" that the respective probability is assigned to. Additionally, in order to calculate a point forecast modes and interquartile ranges we need to construct artificial discrete probability distributions. To do this each point forecast is placed into a bin and the bin midpoint substitutes as the point forecast value. We use inflation bins that are 0.25 percentage points wide and range from -1 to 7 percent (24 bins) and GDP bins that are 0.5 percentage points wide and range from -2 to 8 percent (20 bins).

- vInf: median across forecasters of the standard deviation of probability-based inflation forecasts
- vGDP: median across forecasters of the standard deviation of probability-based inflation forecasts
- DvInf: interquartile range across forecasters of inflation point forecasts
- DvGDP: interquartile range across forecasters of GDP point forecasts
- sInf: median across forecasters of mean less mode of probability-based inflation forecasts
- sGDP: median across forecasters of mean less mode of probability-based GDP forecasts
- DsInf: mean across forecasters of inflation point forecast less inflation point forecast mode
- DsGDP: mean across forecasters of GDP point forecast less GDP point forecast mode

F Test for risk affecting the policy rule coefficients

Without loss of generality, consider the following simplified policy rule in which uncertainty affects the responsiveness of the policy rate R_t to uncertainty about the inflation forecast (notation is not the same as in the main text):

$$R_t = \beta(\sigma_t)\pi_t^f + u_t, \tag{1}$$

where π_t^f denotes the time t inflation forecast, σ_t denotes uncertainty over the inflation forecast, and $\beta(\sigma_t)$ is given by

$$\beta(\sigma_t) = \beta_0 + \beta_1\sigma_t^2.$$

It follows that (1) can be written

$$R_t = \beta_0\pi_t^f + \beta_1\sigma_t^2\pi_t^f + u_t \tag{2}$$

Suppose that instead of estimating (2) one estimates

$$R_t = \gamma_0\pi_t^f + \gamma_1\sigma_t^2 + \tilde{u}_t.$$

It is straightforward to show that the ordinary least squares estimate of $\Gamma = [\gamma_0, \gamma_1]'$ can be expressed as

$$\hat{\Gamma} = \begin{pmatrix} \beta_0 + \beta_1 \dfrac{\Sigma'\Sigma\Pi'\Phi - \Pi'\Sigma\Sigma'\Phi}{\Pi'\Pi\Sigma'\Sigma - (\Pi'\Sigma)^2} \\ \beta_1 \dfrac{\Pi'\Pi\Sigma'\Phi - \Pi'\Sigma\Pi'\Phi}{\Pi'\Pi\Sigma'\Sigma - (\Pi'\Sigma)^2} \end{pmatrix},$$

where Σ is the $T \times 1$ column vector containing the T time series observations on σ_t^2; Π is the $T \times 1$ column vector containing the T time series observations on π_t^f; and Φ the $T \times 1$ column vector containing the T time series observations on $\pi_t^f \sigma_t^2$. It follows that a test of the null hypothesis $\gamma_1 = 0$ corresponds to a test of $\beta_1 (\Pi'\Pi\Sigma'\Phi - \Pi'\Sigma\Pi'\Phi) = 0$. As long as $\Pi'\Pi\Sigma'\Phi - \Pi'\Sigma\Pi'\Phi \neq 0$ (in large samples), then, testing for $\gamma_1 = 0$ is equivalent to a test of $\beta_1 = 0$.

Of course this test will not have any power if $\Pi'\Pi\Sigma'\Phi - \Pi'\Sigma\Pi'\Phi = 0$ (in large samples). In large samples this latter condition is:

$$\operatorname*{plim}_{T\to\infty} \frac{\sum_{t=1}^T \pi_t^{f2}}{T} \frac{\sum_{t=1}^T \pi_t^f \sigma_t^4}{T} - \frac{\sum_{t=1}^T \pi_t^f \sigma_t^2}{T} \frac{\sum_{t=1}^T \pi_t^{f2} \sigma_t^2}{T} = 0.$$

This could occur if π_t^f and σ_t^2 are independent and $E\pi_t^f = 0$ in which case

$$\operatorname{plim} \sum_{T\to\infty} \pi_t^f \sigma_t^4 / T = E\left[\pi_t^f \sigma_t^4\right] = E\pi_t^f E\sigma_t^4 = 0$$

and

$$\operatorname{plim} \sum_{T\to\infty} \pi_t^f \sigma_t^4 / T = E\left[\pi_t^f \sigma_t^4\right] = E\pi_t^f E\sigma_t^4 = 0.$$

References

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

Jurado, K., S. Ludvigson, and S. Ng (2015). Measuring uncertainty. *American Economic Review 105*(3), 1177–1276.

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

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

Thornton, D. L. (2005). A new federal funds rate target series: September 27, 1982 – december 31, 1993. Federal Reserve Bank of St. Louis working paper 2005-032.