Macroeconomic Dynamics, **23**, 2019, 1793–1814. Printed in the United States of America.
doi:10.1017/S1365100517000451

"WAIT-AND-SEE" MONETARY POLICY

XIAOWEN LEI
Simon Fraser University

MICHAEL C. TSENG
Swiss Finance Institute

This paper develops a model of the optimal timing of interest rate changes. With fixed adjustment costs and ongoing uncertainty, changing the interest rate involves the exercise of an option. Optimal policy therefore has a "wait-and-see" component, which can be quantified using option pricing techniques. We show that increased uncertainty makes the central bank more reluctant to change its target interest rate, and argue that this helps explain recent observed deviations from the Taylor Rule. An optimal wait-and-see policy fits the target interest rates of the Fed and Bank of Canada better than the Taylor Rule.

Keywords: Monetary Policy, Interest Rate Inertia, Optimal Stopping

One possible FOMC strategy is to simply pocket the lower yields and continue to wait-and-see on the U.S. economic outlook. James Bullard, Federal Reserve Bank of St. Louis President, June 5, 2012 Speech.

1. INTRODUCTION

Interest rate adjustments occur at a lower frequency than macroeconomic data releases. For example, Figure 1(a) plots the monthly US federal funds target rate against the inflation rate. The frequency mismatch is apparent. Of course, in recent years interest rates have remained at zero for standard liquidity trap reasons, but rates exhibited inertia even before the financial crisis. Such inertia is not specific to the US Federal Reserve. The same plot for the Bank of Canada (BoC), Figure 1(b), shows that the BoC rates exhibit even more pronounced inertia.

There are several possible reasons why a central bank might be reluctant to change interest rates. The most obvious reason is that there is some sort of cost

We gratefully acknowledge Kenneth Kasa, without whose guidance, patience, and encouragement this paper would not have been possible. We are also especially grateful to an anonymous referee for many useful suggestions. We also thank Galo Nuño, James Costain, Janet Hua Jiang, Jinill Kim, Edouard Djeutem, David Andolfatto, Robert Jones, Luba Petersen, John Knowles, Lucas Herrenbrueck, Peter Zadrozny, and participants in the Simon Fraser University Economics PhD seminar, and Simon Fraser University Economics Brown Bag seminar, for helpful comments. We are responsible for all the remaining errors. Address correspondence to: Xiaowen Lei, Department of Economics, Simon Fraser University, WMC 2700, 8888 University Dr, Burnaby, BC, V5A 1S6, Canada; e-mail: xiaowen_lei@sfu.ca.

 1365-1005/17

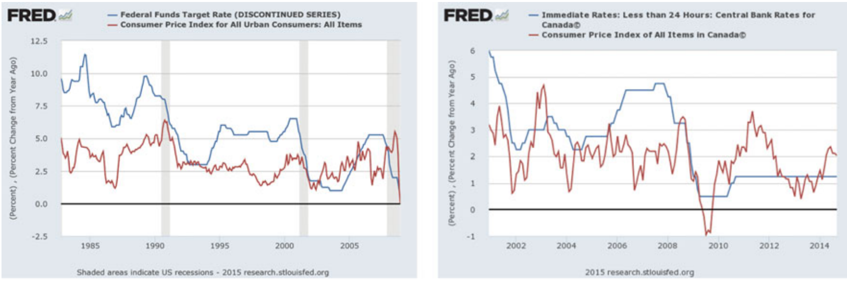

(a) (b)

FIGURE 1. US federal fund rate and Bank of Canada annual interest rate vs. inflation, monthly data. (a) US federal fund rate vs. inflation. (b) Bank of Canada annual interest rate vs. inflation.

to changing the interest rate. However, standard convex adjustment cost models would not explain the observed long periods without change. Such models instead generate *continuous* gradual adjustment. Another possibility is that inertia arises for strategic reasons [see, e.g., Woodford (1999)]. Once again, however, strategic inertia produces continuous adjustment. A third possibility is insufficient data, as most national accounts data arrive only quarterly. Two facts run counter to this explanation. First, highly relevant financial market data are now available at virtually a continuous rate. Second, the plots in Figure 1 show that central banks respond even more haltingly than the frequency at which data are released.

Another information-based explanation comes from the demand-side rather than the supply-side. Perhaps, the Fed simply cannot *process* all the information that is available. Therefore, "rational inattention" produces inertia. This is also contradicted by the facts. First, standard linear-quadratic models of rational inattention produce attenuated adjustments, not discrete adjustments [see Sims (2010)]. Second, the Fed employs *thousands* of economists and data analysts, whose primary job is to process data. Although rational inattention might be a plausible explanation for a lone central banker, it seems less plausible when applied to an entire institution of data-processing specialists. Finally, probably the most common explanation is that wait-and-see is motivated by learning, i.e., the desire to reduce uncertainty. Despite its apparent plausibility, keep in mind standard learning models produce continuous adjustment, or at least adjustment at the same frequency at which the data are available. Even more problematic is the prediction that learning should generate trends in the data. As learning gradually reduces uncertainty, the central bank should become more willing to respond. Such trends are not apparent in the data.

Although all of these factors are no doubt important parts of monetary policy, they do *not* explain infrequent adjustment. Instead, in our view such inaction signals the presence of a *fixed cost* to taking action. What is this cost? Clearly, it is not a technical cost. The Fed could dictate a minute-by-minute target if it so desired. One possibility arises from the fact that monetary policy decisions are not made by a single individual, but by a *committee*. As anyone who has ever served on a committee can attest, committee decision-making has costs. We focus on the inertia induced by these costs. It is possible that committee decision-making introduces inertia by itself, independent of option value considerations, which arise from possible delays in reaching consensus.[1] However, Blinder (2009) cites evidence that it does not. Moreover, Blinder (2009) argues that such procedural delays are likely to be relatively low in the case of US and Canadian monetary policy, at least for the time period studied here, since they are both examples of what he calls "autocratically collegial" committees, in which a single chairman has a dominant influence.

The Federal Open Market Committee (FOMC) must meet at least four times each year in Washington, D.C. Since 1981, eight regularly scheduled meetings have been held each year at intervals of five to eight weeks. At each regularly scheduled meeting, the committee votes on the policy to be carried out during the interval between meetings. Other central banks have similar meeting schedules. Although in theory the Fed can change rates between meetings, this rarely happens. According to FOMC statements, since 2002, 47 out of 52 target rate changes occurred at meeting times. Moreover, as the above plots reveal, it is not uncommon for rates to remain unchanged following a meeting of the FOMC. This suggests that there is more to infrequent adjustment than the costs of holding meetings.

In this paper, we take the presence of fixed costs as given, and study their implications in a general equilibrium framework. We consider a continuous-time version of the standard New Keynesian model, in which the central bank attempts to balance inflation and output gap deviations. Without fixed costs, optimal policy would produce a version of the Taylor Rule. Continuous evolution of macroeconomic data would produce continuous interest rate adjustments. Our primary goal is to show that the option value of waiting-to-see does not just explain infrequent adjustment. If Taylor Rule predictions were simply wiggles around the implied wait-and-see policy, the value added here would be second order. Instead, we argue that an option value perspective introduces fundamentally new and important considerations into policy, which help to explain several widely noted and debated discrepancies between observed interest rates and Taylor Rule predictions. In particular, we show that *uncertainty* becomes an important input into policy. In contrast, standard Taylor Rule models are based on certainty-equivalence, and therefore predict that uncertainty does not influence policy.[2] Greater uncertainty increases the option value of waiting. We argue that post-2001 in the United States is a period with higher uncertainty than pre-2001, thus policy makers react with more inertia in interest rate changes. In calibrating the model to US data, we find that the implied fixed cost corresponds to an approximately 1.89% of additional annual welfare loss.

Before proceeding, we note some important caveats to our analysis. First, although the novelty here is to incorporate option value considerations into macroeconomic policy, these option values are no doubt present in the private sector as well. In fact, studying these problems is where the analytical techniques were first developed [Dixit (1994), Stokey (2008)]. For simplicity, we assume that only the policy maker faces fixed costs. There are reasons to believe, however, that waiting-to-see could interact in mutually destructive ways if both the government and the private sector have an option value to wait-and-see, especially if each side lacks information about the other [Caplin and Leahy (1994)]. Second, the only uncertainty in our model is that regarding the future values of exogenous shocks. The central bank is fully informed about the underlying structural model and its parameter values. As noted above, we doubt whether learning about parameters would, by itself, explain observed inertia. However, if policy makers confront more diffusive forms of uncertainty, in combination with fixed costs, this might influence the implied option value of waiting-to-see [see, e.g., Miao and Wang, Neng (2011) and Trojanowska and Kort (2010)].

The remainder of the paper is organized as follows. Section 2 relates our work on optimal monetary policy to the real options literature and to recent debates about Taylor Rule deviations. Section 3 derives a benchmark model without fixed costs of adjustment, which is essentially a version of the Taylor Rule. Section 4 solves for partial equilibrium, where the private sector's expectation is exogenous. Section 5 proves that the central bank's optimal policy is part of a rational expectation equilibrium. Section 6 examines the properties of an optimal "wait-and-see" rule with respect to parameter changes. Section 7 takes the model's testable predictions to both US Fed and BoC data and studies its quantitative implications. Section 8 discusses future research in this framework and concludes. The appendix contains a technical proof for sufficiency of appropriate Euler equations for optimality in terms of general stopping time problems.

2. LITERATURE REVIEW

Our study is motivated by recent debates about Taylor Rule deviations. Kahn (2010) examines various versions of the Taylor Rule, and shows that although a Taylor Rule matches federal funds target rate data well during 1980–1990s, there are large deviations from it in recent years. In particular, policy appears to have been too accommodative. Taylor (2014) argues that interest rate policy has become less predictable, and that the central bank has been "keeping rates too low for too long." In addition, Woodford (1999) shows that the optimal nominal interest rate rule should exhibit "inertia." In our model, the central bank exhibits a type of inertia different from that of Woodford. Here, higher inertia means the central bank widens the "no-action" band in which it allows the state to drift freely without applying interest rate adjustment. In contrast, Woodford shows that the ability to precommit to an interest rate path introduces an autoregressive component into the Taylor Rule. A larger autoregressive component corresponds to higher inertia. The idea is that small adjustments in the same direction improve the intertemporal trade-off between inflation and the output-gap [Rudebusch (2006)]. However, in practice, interest rates exhibit not only small adjustment but also large shifts, and, more recently, periods of no adjustment. Such time series does not have the typical sample path properties of an autoregressive process, suggesting possible misspecification. Rather than tying the current interest rate to the last adjustment via an autoregressive mechanism, in our model the central bank decides whether to adjust based on the current state of the economy. In this "wait-and-see" approach, apparent inertia arises from an option value. Periods of no adjustment and occasional large shifts, unaccounted for in an autoregressive specification, are naturally explained in our model by the central bank's reluctance to exercise an option.

In related work, Davig and Leeper (2008) and Svensson and Williams (2008) compute optimal Markov switching rules in which a policy maker sets different degrees of response to inflation depending on whether the state has crossed an *exogenous* threshold. Alba and Wang (2017) and Murray et al. (2015) fit empirical Markov switching models to US monetary policy. In these models, discrete adjustment reflects a switch between two sets of Taylor Rules. In contrast, we study optimal policy with endogenously set optimal thresholds. Our model is also related to prior work on "sticky information." For example, Mankiw and Reis (2002) study the impact of discrete, optimally chosen, information updating on the Phillips curve. However, their focus is on the firm's side.

There has been considerable effort in bringing the real option effect into macroeconomic models [see, e.g., Froot and Obstfeld (1989), Dixit (1993), Dixit (1994), and Stokey (2008)], which drives the key "wait-and-see" feature of our model. The techniques that we employ are also used in the menu cost literature. When firms face both first and second moment uncertainty about total-factor productivity, Bloom et al. (2012) show that uncertainty increases during recessions. The nonconvexities together with time variation in uncertainty imply that firms become more cautious in investing when uncertainty increases. Stokey (2016) develops a model which shows that uncertainty about future tax policy induces firms to temporarily stop investing. Both papers look at uncertainty from the firm side. A recent paper that examines uncertainty from the policy side is Alvarez and Dixit (2014), which analyzes the optimal timing of a break-up of the Euro zone using a real options framework, with the private sector modeled in a reduced form way. In contrast, our model casts a standard New Keynesian model into an impulse control framework and remains in a general equilibrium setting where both the central bank and private section have rational expectations.

3. MONETARY POLICY WITHOUT FIXED COSTS

We briefly recall the optimal monetary policy under discretion in the absence of fixed costs. Assume time is discrete. Let π_t denote inflation and x_t denote the output gap. The Phillips curve and the dynamic investment/saving (IS) curve in a standard New Keynesian model, are

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

and

$$x_t = -\frac{1}{\gamma}[i_t - E_t(\pi_{t+1}) - r] + E_t(x_{t+1}), \tag{2}$$

where γ is the household's coefficient of risk-aversion, κ denotes the response of inflation to an output shock, i_t is the nominal interest rate, and r is the natural rate of interest, which is assumed to be constant here. The cost-push shock u_t is essential in our analysis. It represents the uncertainty the central banker is facing.

The central banker in our model conducts monetary policy with discretion, and therefore has the myopic goal of minimizing the deviations of current inflation and the output gap from target (normalized to zero).[3] Without adjustment costs, the central banker's problem is a static one:

$$\min_{x_t,\pi_t} \frac{1}{2}(\lambda x_t^2 + \pi_t^2), \tag{3}$$

subject to

$$\pi_t = \kappa x_t + v_t, \tag{4}$$

where $v_t \equiv \beta E_t(\pi_{t+1}) + u_t$, and λ is the weight put on the output gap.[4] Assuming u_t is an AR(1) process with coefficient ρ_u, the optimal interest rate rule is given by

$$i_t^* = r + \Phi_i u_t, \tag{5}$$

where $\Phi_i = \frac{\kappa\gamma(1-\rho_u)+\lambda\rho_u}{\kappa^2+\lambda(1-\beta\rho_u)}$. In the special case where u_t follows a random walk, we have $\Phi_i = \frac{\lambda}{\kappa^2+\lambda(1-\beta)}$. It is well known that, without fixed costs, a linear-quadratic objective function implies certainty equivalence for the central bank's problem. *Uncertainty plays no role in optimal monetary policy, and only the current state matters*. As expected, the optimal interest rate is linear and *continuous* in the cost push shock. Under discretion, the output gap and inflation are also linear in the cost-push shock. Thus, we can view this linear rule as a version of the Taylor Rule, which will be compared with the central bank's optimal rule from our model in Section 4.

4. MODEL

We continue to assume the bank operates with discretion. The model is obtained by taking continuous time limits of the previous two standard discrete time New Keynesian equations [see, e.g., Galí (2009)]. Instead of solving a continuous-time New Keynesian model from the outset [as is done, for example, in Fernández-Villaverde et al. (2012)], we start by deriving discrete time counterpart, which summarizes the private sector's behavior under any policy, then take the continuous time limit. Our New Keynesian continuous-time model admits analytic solutions, using techniques from stochastic control.

Rational expectations require the private sector to form expectations of the output gap and inflation that are consistent with central bank policy. Given such expectations, the central bank chooses optimal monetary policy taking the private sector's expectations as given. The resulting optimal policy must then conform to the private sector's expectations. In discrete time, it is reasonable to conjecture that the private sector's expectations follow a martingale[5]:

$$E_t \pi_{t+1} = \pi_t, \tag{6}$$

$$E_t x_{t+1} = x_t. \tag{7}$$

In a discrete time setting, the private sector's information set at time t includes the current period interest rate. In our continuous-time model, the interest rate is determined by the current state of the economy, which is evolving stochastically. The kth adjustment i_k happens at a random, rather than a deterministic, time τ_k. Therefore, the continuous-time counterparts to the above conditions for stochastic processes π_t and x_t are that, conditional on the central bank's last adjustment, the private sector's expectation remains the same between interest rate adjustments:

$$E[\pi_\tau | \tau', i_k] = \pi_{\tau'} \quad \text{and} \quad E[x_\tau | \tau', i_k] = x_{\tau'},$$

where τ and τ' are any stopping times with $\tau_{k+1} > \tau \geq \tau' \geq \tau_k$.[6]

Substituting (6) and (7) into (1) and (2), respectively, gives

$$x_t = a_1 \tilde{i} + b_1 u_t, \tag{8}$$

$$\pi_t = a_2 \tilde{i} + b_2 u_t, \tag{9}$$

where $\tilde{i} = i_k - r$ for some k. The cost-push shock u_t is exogenous and is modeled by a Brownian motion.[7]

We derive a rational expectations equilibrium in which the central bank incorporates ex-ante assumptions about the private sector's expectations into his cost minimization problem. Then, variables x_t and π_t in the objective function in the central bank's problem can be summarized by a single state variable:

$$\begin{aligned}
\frac{1}{2}(\lambda x_t^2 + \pi_t^2) &= \frac{1}{2}[\lambda(a_1 \tilde{i} + b_1 u_t)^2 + (a_2 \tilde{i} + b_2 u_t)^2] \\
&= \frac{1}{2}(\lambda a_1^2 + a_2^2)\tilde{i}^2 + (\lambda b_1^2 + b_2^2)u_t^2 + 2(a_1 b_1 \lambda + a_2 b_2)\tilde{i} u_t \\
&= \frac{1}{2}\tilde{\lambda}(\tilde{i} - \theta u_t)^2 + \frac{\lambda^2 \beta(2-\beta)}{[\lambda(1-\beta)^2 + \kappa^2]^2} u_t^2,
\end{aligned} \tag{10}$$

where $\theta = \frac{\lambda(1-\beta)}{\lambda(1-\beta)^2 + \kappa^2}$, and $\tilde{\lambda} = \lambda a_1^2 + 1$.

Define $z_t \equiv \tilde{i}_t - \theta u_t$. Note that the second term in (10) is independent of monetary policy. It is the efficiency loss that cannot be mitigated when adjusting the interest rate. The first term implies a squared deviation of optimal interest rate. Assuming the above assumption holds, we could derive the conjectured coefficients $a_1 = \frac{\beta - 1}{\kappa}$, $b_1 = -\frac{1}{\kappa}$, $a_2 = -1$, $b_2 = 0$.

Let K denote the fixed cost of adjustment. A discretionary policy maker's problem is to minimize the expected discounted sum of squared deviations of z_t and adjustment cost by choosing a sequence $\{\tau_k, i_k\}$, where τ_k and i_k are the time and size, respectively, of the kth adjustment. The central bank's problem is then to find

$$V(z_0) = \inf_{\tau_k, i_k} E^{z_0}\left[\int_0^{\infty} e^{-\rho t}\left[\frac{1}{2}\tilde{\lambda} z_t^2\right] dt + \sum_{k=1}^{\infty} e^{-\tau_k} K\right], \quad \textbf{(11)}$$

where

- The expectation operator $E^{z_0}[\cdot]$ denotes expectation taken with respect to the law of the shock process z_t.
- The cost-push shock u_t is assumed to follow a Brownian motion $du_t = \sigma dB_t$, where σ is the volatility of the process. Increasing σ increases the uncertainty faced by the central bank.
- From (10), the summarized variable z_t then evolves according to $dz_t = -\theta\sigma dB_t = du_t$ with initial state z_0 almost surely.
- $\{\tau_k\}$ is a nondecreasing sequence of stopping times with respect to the natural filtration $\{\mathcal{F}_t\}$ generated by du_t. The timing of the kth adjustment, which is a state-dependent random time τ_k, must be decided only using the central bank's information.
- Each i_k, the size of kth adjustment, is a τ_k-measurable random variable. In other words, the central bank determines the size of an adjustment using information available at the (state-dependent random) time of adjustment.
- $z_{\tau_k} - z_{\tau_k^-} = i_{\tau_k} - i_{\tau_k^-}$. Since the cost-push shock is exogenous, changes in z_t correspond to changes in the interest rate.

Therefore, the central bank's problem is to choose a random state-dependent sequence of adjustment times based on information available, along with the interest rate at those times. In the presence of a fixed cost K, continuous adjustment is clearly suboptimal.

The central bank's problem can be seen as a dynamic programming problem. Define the "best adjustment operator" $\mathcal{A}$, acting on bounded functions $W : \mathbf{R} \to \mathbf{R}$, by[8]

$$\mathcal{A}W(z) = \inf_{i \in \mathbf{R}} W(z - i) + K.$$

If $V(z)$ is the central bank's value function, then $\mathcal{A}V(z)$ is the resulting value from the best possible interest rate adjustment, if an adjustment is made at state z. Therefore, we must have

$$V(z) = \inf_{\tau} E^z\left[\int_0^{\tau} e^{-\rho t} f(z_t) d_t + e^{-\rho\tau}\mathcal{A}V(z_{\tau^-})\right],$$

where, for ease of notation, we put $f(z) = \frac{1}{2}\tilde{\lambda} z^2$ and $\inf_\tau$ denotes the infimum over all finite stopping times. By optimality, $V(z) \le \mathcal{A}V(z)$ since the central bank always has the choice of not exercising the option of adjusting. The central bank's problem can be therefore further rewritten as follows:

$$V(z) = \inf_{\tau} E^z \left[\int_0^{\tau} e^{-\rho t} f(z_t) d_t + e^{-\rho \tau} V(z_{\tau^-}) \right].$$

Assuming now that $V(z) \in C^2(\mathbf{R})$, i.e., is twice differentiable with continuous second derivative and the infimum $\inf_\tau$ is actually attained at a stopping rule τ^* given by an open region $U \subset \mathbf{R}$, then applying Dynkin's Formula to the right-hand side gives[9,10]

$$V(z) = V(z_0) + E^z \left[\int_0^{\tau^*} e^{-\rho t} \left[f(z_t) - \rho V(z_t) + \frac{1}{2} V''(z_t) \theta^2 \sigma^2 \right] dt \right]. \quad \mathbf{(12)}$$

Therefore, the state space $\mathbf{R}$ is divided into two regions

$$\begin{cases} \rho V = \frac{1}{2} \tilde{\lambda} z^2 + \frac{1}{2} V'' \theta^2 \sigma^2 & \text{on } U \\ \rho V < \frac{1}{2} \tilde{\lambda} z^2 + \frac{1}{2} V'' \theta^2 \sigma^2 & \text{on } \mathbf{R} \backslash U. \end{cases}$$

The inaction region U is therefore characterized by the central bank's first-order condition[11]

$$\rho V = \frac{1}{2} \tilde{\lambda} z^2 + \frac{1}{2} V'' \theta^2 \sigma^2, \quad \mathbf{(13)}$$

between adjustment and nonadjustment of the interest rate. It states that the expected loss of a central banker's value is equal to the sum of an immediate squared deviation of optimal interest rate, and the expected rate of capital loss of holding the option of not changing rates. Adjustment is applied immediately on $\mathbf{R} \backslash U$.

Conjecture now that U is an interval $U = (a, b)$ for some $a < b$. On U, the first-order condition is an inhomogeneous ordinary differential equation with general solution

$$V = A z^2 + B_1 e^{\delta_1 z} + B_2 e^{\delta_2 z} + C,$$

with unknown constants A, B, δ_1, δ_2, and C. The constants A, B, δ_1, δ_2 are found by matching coefficients as follows:

$$A = \frac{\tilde{\lambda}}{2\rho}, C = \frac{\tilde{\lambda} \theta^2 \sigma^2}{2\rho^2}, \delta_{1,2} = \pm \frac{\sqrt{2\rho}}{\theta \sigma}. \quad \mathbf{(14)}$$

What remains is to solve for the free boundaries a and b, and to determine the adjustment amount i once the boundary is reached. We note that that $z - i = c$ must lie in (a, b). Furthermore, to satisfy the C^2 assumption on V in the above derivation, it is necessary that $V(a) = V(b) = V(c) - K$, $V'(a) = V'(b) = 0$, and $V''(a) = V''(b) = 0$. We prove in the appendix that the first two set of conditions

$$V(a) = V(b) = V(c) - K, \quad \mathbf{(15)}$$

$$V'(a) = V'(b) = 0, \quad \mathbf{(16)}$$

are sufficient for V to be the value function by generalizing Dynkin's Formula beyond C^2-functions. In other words, the optimizing central bank needs only to satisfy continuity, (15), and an Euler equation, (16), on the adjustment thresholds. The value function V is constant $V = V(c)$ on $\mathbf{R}\backslash(\mathbf{a}, \mathbf{b})$. When facing a starting state $z_0 \in \mathbf{R}\backslash(\mathbf{a}, \mathbf{b})$, the central bank immediately adjusts the state to c. Given the model's symmetry, it is reasonable to guess that (a, b) is of the form $(-S, S)$ and $c = 0$. Using this guess, the Euler equations become

$$AS^2 + B_1 e^{\delta_1 S} + B_2 e^{\delta_2 S} = B_1 + B_2 - K, \tag{17}$$

$$2AS + B_1\delta_1 e^{\delta_1 S} + B_2\delta_2 e^{\delta_2 S} = 0, \tag{18}$$

respectively. This is the same set of equations encountered in Dixit (1993). An approximate analytical solution for the threshold S is

$$S^* \approx \left(\frac{12\theta^2\sigma^2 K}{\tilde{\lambda}}\right)^{1/4}. \tag{19}$$

This implies that a fixed cost K of fourth-order magnitude has first-order importance in determining the adjustment threshold, while uncertainty of second order has first-order effects on inertia.

5. RATIONAL EXPECTATIONS

It is clear that the inflation process π_t and the output process x_t, under the central bank's optimal interest rate policy, are consistent with the private sector's expectations. Indeed, recall that the private sector's expectations regarding the two processes satisfy a local martingale-type condition:

$$E[\pi_\tau|\tau', i_k] = \pi_{\tau'} \quad \text{and} \quad E[x_\tau|\tau', i_k] = x_{\tau'},$$

where τ and τ' are any stopping times with $\tau_{k+1} > \tau \geq \tau' \geq \tau_k$. On the other hand, under the central bank's optimal policy,

$$x_\tau = a_1 i_k + b_1 u_\tau,\ x_{\tau'} = a_1 i_k + b_1 u_{\tau'}$$

and

$$\pi_\tau = a_2 i_k + b_2 u_\tau,\ \pi_\tau = a_2 i_k + b_2 u_{\tau'}.$$

By Doob's Optional Stopping Theorem [Lipster and Shiryayev (1989)], we have

$$E[x_\tau|\tau', i_k] = a_1 i_k + b_1 u_{\tau'} = x_{\tau'}$$

and

$$E[\pi_\tau|\tau', i_k] = a_2 i_k + b_2 u_{\tau'} = \pi_{\tau'}.$$

Hence, expectations are confirmed.

6. EXPECTED HITTING TIMES AND ADJUSTMENT

The properties of the optimal rule in our model shed some light on how different uncertainty regimes contribute to monetary policy inertia.[12] Given an arbitrary initial point z, the expected time to the first such resetting is found by[13]

$$T(z) = \frac{S^2 - z^2}{\sigma^2}. \tag{20}$$

Expressed in terms of the endogenous threshold $\pm S^*$ from equation (18), expected hitting time is

$$T(z) = \frac{\sqrt{\frac{12\theta^2 K}{\tilde{\lambda}}}}{\sigma} - \frac{z^2}{\sigma^2}. \tag{21}$$

The derivative of $T(z)$ with respect to σ is

$$\frac{\partial T(z)}{\partial \sigma} = \frac{2(z^2 - \sigma\sqrt{3\theta^2 K/\tilde{\lambda}})}{\sigma^3}. \tag{22}$$

Thus, the effect of uncertainty on expected hitting times depends on the current state of the economy. There are two countervailing effects. On the one hand, given a fixed threshold, a higher variance increases the probability that z hits the boundary, so the expected hitting time decreases. On the other hand, increased uncertainty also increases the optimal threshold, which makes the state less likely to reach the threshold. In our case, when $z^2 < \sigma\sqrt{3\theta^2 K/\tilde{\lambda}}$, the first effect dominates, and increased uncertainty decreases expected hitting times. This corresponds to the case where the initial interest rate is nearly optimal, but if the economy is hit with a severe negative shock, the central banker must react by immediately lowering interest rate.[14] If $z^2 > \sigma\sqrt{3\theta^2 K/\tilde{\lambda}}$, the second effect dominates, and increased uncertainty would increase expected hitting time. This could be the scenario where, given a mild level of inflation, the central banker would have increased the interest rate during normal times, but is less likely to do so during a recession. The two plots in Figure 1 show that both effects are present in the data. A large downward adjustment usually corresponds to an unexpected recession followed by a longer period of interest rate inertia. For example, consider the two recession periods of the US economy during early 90's and early 2000's. In both the cases, the Fed lowered target rates swiftly at first, then waited for a long period (15 months and 9 months, respectively) before further adjustments.

7. CALIBRATION

Our model has testable implications for observed interest rate targets. In particular, we take our wait-and-see rule to US Fed and BoC interest rate data, and examine whether the model can explain both the observed inertia and the timing of shifts in rates. For the United States, monthly industrial production, inflation, and federal funds rate data for the period October 1982 to July 2008 are obtained from the

TABLE 1. Parametrization (monthly)

Description	Notation	Value
Risk aversion	γ	0.5
Discount rate	β	0.9967
Natural rate (United States)	r_{USA}	0.2385%
Natural rate (Canada)	r_{Canada}	0.1570%
Sensitivity of inflation on output	κ	0.0076
Weight on output deviation	λ	1/2
Adjustment cost (Fed)	K_{Fed}	0.0126
Adjustment cost (BoC)	K_{BoC}	0.1300
United States before 2001 volatility	σ_l	0.0126%
United States after 2001 volatility	σ_h	0.0242%
Canada volatility	σ	0.0135%

Federal Reserve Bank of St. Louis website. Since monthly data on GDP are not available, industrial production serves as a proxy for the output variable.[15] To compute potential monthly output, a Hodrick–Prescott filter with $\lambda = 14{,}400$ is used.

Using the parameters reported in Table 1, we calibrate the model to the observed federal funds target. The initial model interest rate is calibrated to match the actual initial interest rate in the sample period. The risk aversion parameter γ is set to be 0.5. The monthly discount factor β is 0.9967, to match an average annual discount rate of 4%. We calibrate the natural interest rate r from the sample average estimate in Holston et al. (2017) of United States from 1990 to 2007. A similar sample average estimate is computed for Canada from 2007 to 2015. This gives $r_{\text{USA}} = 2.9\%$ yearly, implying $r_{\text{USA}} = 0.2385\%$ monthly, and $r_{\text{Canada}} = 1.9\%$ yearly, implying $r_{\text{Canada}} = 0.1570\%$ monthly. The weight on output deviation is 0.5. We take monthly $\kappa = 0.0076$ from Galí and Gertler (1999).[16] We then break the US data into two periods, before and after 2001, since we view the post-2001 period as being more volatile (e.g., the tech stock boom and bust, the 9/11 attack, the financial crisis). The parameters σ_l and σ_h are then computed from the standard deviation of the u_t series in these two periods. The u_t process is inferred from output gap and inflation data, along with the Phillips curve and model-implied private sector beliefs.[17,18] To be more concrete, we know that the New Keynesian Phillips curve gives us[19]

$$\pi_t = \kappa x_t + \beta \mathbf{E}_t(\pi_{t+1}) + u_t. \tag{23}$$

We also know from the model that $\mathbf{E_t}(\pi_{\mathbf{t+1}}) = \pi_{\mathbf{t}}$. Therefore, we can infer the following cost shock process:

$$u_t = (1 - \beta)\pi_t - \kappa x_t. \tag{24}$$

TABLE 2. Mean-squared error, percent

Central banks	Model	Taylor
US Fed	2.04	2.17
BoC	1.11	1.20

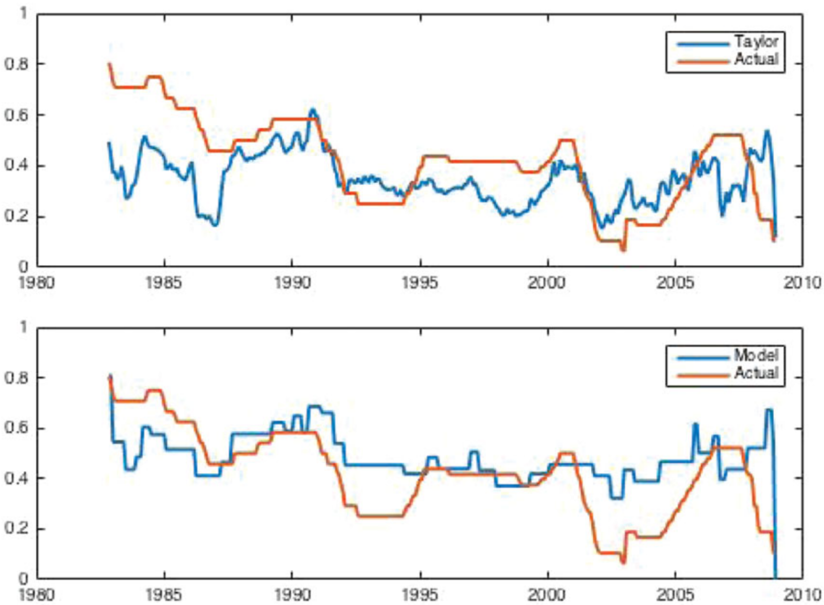

FIGURE 2. US: Model vs. Taylor Rule.

We then calibrate the adjustment cost K to match the weighted average interest rate duration (weighted by its duration). With US data, this produces $K_{\text{Fed}} = 0.0126$. With Canadian data, this produces $K_{\text{BoC}} = 0.13$. We will explain the quantitative interpretation of this shortly. Figures 2 and 3 demonstrate how the wait-and-see rule compares with the Taylor Rule. The main visual difference is that our rule exhibits policy inertia and large discrete jumps.

The adjustment cost parameters K_{Fed} or K_{BoC} can be translated to loss function equivalents. $K = 0.0126$ implies a fixed cost of 1.26% standard deviation from the steady state. The median interest rate duration in the United States is 8 months, which implies the Fed changes rates about 1.5 times per year on average. Therefore, the annual cost is approximately an extra $1.5\% \times 1.26\% = 1.89\%$ of welfare loss. Using the above parameters, we compare the goodness of fit of the Taylor Rule and the wait-and-see rule: The result is shown in Table 2. The result that

TABLE 3. Probability of leaving rates unchanged

Central banks	Actual	Model	Taylor
US Fed	0.5	0.8949	0
BoC	0.7410	0.9398	0

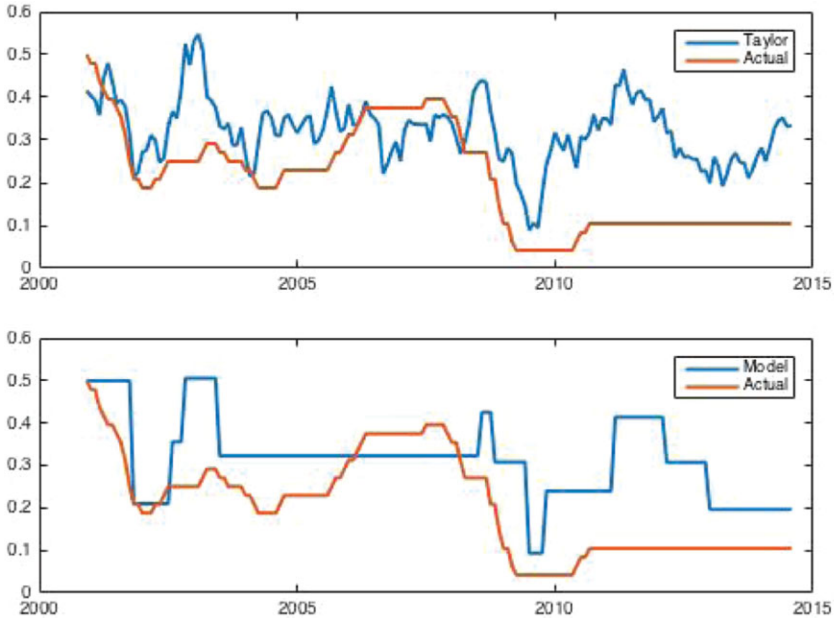

FIGURE 3. Canada: Model vs. Taylor Rule.

wait-and-see rule has a better overall fitness is also robust to alternative parametrization of discount rate and elasticity of inflation to output parameter.

The previous results calibrate the adjustment cost K to match the weighted average interest rate duration, and study the overall fitness. For a robustness check, one can also calibrate K to minimize the mean-squared error between the target rate and the model-implied interest rate, and study how well they match the probability of interest rate change and the average size of interest rate changes. This produces $K_{\text{Fed}} = 0.017$, and $K_{\text{BoC}} = 0.1$ when keeping all other parameters with the same value as Table 1. Tables 3–5 report the results.

We compare three statistics associated with the policy rule from our model and the Taylor Rule.

(a) Probability of nonadjustment. Table 3 displays an obvious distinction between the Taylor Rule and the wait-and-see rule. The Taylor Rule produces a

TABLE 4. Conditional average absolute percentage change of rates

Central banks	Actual	Model	Taylor
US Fed	0.0531	0.1981	0.0676
BoC	0.1327	0.5843	0.1282

TABLE 5. Unconditional average absolute percentage change of rates

Central banks	Actual	Model	Taylor
US Fed	0.0304	0.0208	0.0676
BoC	0.0344	0.0352	0.1282

zero probability of keeping rates constant, while our model generates a reasonable probability of inaction. While it captures the possibility of not changing rates, which the Taylor Rule cannot explain, the wait-and-see rule generates an inaction probability that is larger than the data.

(b) Conditional average absolute percentage change of rates. Our model also has implications for the magnitude of rate changes. To see this, we take the absolute value of percentage change of rates conditional on changes occurring. The absolute value is taken both due to the symmetry of our model, and because the direction of change is not of immediate interest. We find that the wait-and-see rule generates larger changes than in the data.

(c) Unconditional average absolute percentage change of rates. An arguably better measure of overall fit is to consider the conjunction of the above two statistics. That is, taking into the account both the probability of adjustment and the magnitude of adjustment, we compute the unconditional average absolute percentage change of rates. The last column in Tables 4 and 5 are the same, since the probability of changing rates in the Taylor Rule is always one. However, the wait-and-see rule outperforms the Taylor Rule in both the United States and Canada. One can see that although the Taylor Rule matches the magnitude of changes fairly well, its inability to capture the timing of rate shifts forces it to produce a much larger deviation from the actual interest rate than the wait-and-see rule.

8. CONCLUSION

This is the first paper to study the optimal *timing* of interest rate changes. By combining uncertainty and adjustment costs, we are able to rationalize observed interest rate jumps and inertia using a simple impulse control model. The model is consistent with observed target interest rates of both the Fed and the BoC. By focusing on the question of *when* to change interest rates as opposed to *how much*, our framework sheds light on the essential role that uncertainty plays in policy. This motivates many possible directions for future research.

An immediate economic question is the welfare implications of a wait-and-see monetary policy. Without adjustment costs, a discretionary central bank is able to implement the first best outcome by minimizing output gap and inflation deviations instantaneously. With adjustment costs, it is clear that output and inflation deviations will persist, which generates potentially large efficiency losses.

Our technical framework admits several modifications and generalizations. Rather than one discrete adjustment at a time, the central bank may switch between regimes of continuous adjustment and no adjustment. The BoC time series in Figure 1(b) could very well reflect such a policy. To better match the level of interest rates shown by data, the exogenous policy-independent cost-push shock could be modeled by a more general Lévy process rather than Brownian motion.[20] One could also think about imposing a zero lower bound on interest rates. All such extensions fall under the umbrella of general optimal stopping/option exercise problems.

Of particular relevance is the question of how a zero lower bound would affect the central bank's policy in our model. A binding zero lower bound constrains the central bank's ability to react to shocks, since the only option would be an upward adjustment at the zero lower bound. The central bank might alter its behavior in response to this restriction. Near the zero lower bound, the inability, once the bound is reached, to set a negative nominal rate makes the central banker more sensitive to small negative shocks. The option of a large negative adjustment in response to the next large negative shock is no longer available. As a preventive measure, the central bank would apply downward adjustments more frequently near the bound. Away from the bound, the potential loss that could occur at the bound due to restricted ability to adjust could lead the central banker to make positive adjustments larger than she would in our model when positive shocks occur.

One could also argue that the central bank anticipates upturns and downturns of the economy in setting monetary policy, rather than just reacting to changes in uncertainty levels ex post. This could be addressed by a two-factor stochastic volatility model, where volatility of the cost-push shock follows a stochastic process, in the option exercise framework. The driving factor could be observable or latent. A latent stochastic volatility factor is often argued to be a major concern for monetary policy makers.

Also to be explored is the relationship between learning and a wait-and-see policy. For example, if the natural interest rate is unobservable, waiting not only brings the benefit of more information about the underlying shocks, but also refines estimates of parameter values. One could also envision a model where the private sector also faces fixed investment costs. Firms will delay investing due to policy uncertainty, which then feeds back to more uncertainty by the policy maker. This kind of feedback loop might be an impediment to recovery, with both the central bank and the firm waiting-to-see about each other. Last, future research could dig deeper into the way uncertainty is modeled here. For example, with fears of model misspecification, the central bank might want to react *more* aggressively with respect to the current state. So, it might be important to examine how a preference for robustness would interact with the wait-and-see policy generated from our baseline model.

NOTES

1. We thank an anonymous referee for pointing this out.

2. In fairness, some have argued that the Taylor rule is desirable not because it is optimal within the context of a given model, but rather because it is *robust* to the presence of model uncertainty [see, e.g., Rotemberg and Woodford (1999) and Levin and Williams (2003)].

3. Such preference is a second-order approximation of a constant relative risk aversion (CRRA) household's utility function, which is proved in Rotemberg and Woodford (1999) and Woodford (2003).

4. With discretion, the central bank cannot influence the private sector's expectation in a systematic way. Thus, he treats expectations as an exogenous process when conducting policy.

5. In general, we are not able to tell whether such expectations are unique. In fact, when banks' interest rate threshold is fixed, the private sector could have highly nonlinear expectations, as demonstrated by Davig and Leeper (2008). However, since we focus on an endogenous interest rate band rather than endogenous expectation, studying a version of martingale expectation and verifying that it is rational simplifies the analysis and helps us to focus on the model's inertia predictions.

6. Recall that a *filtration* is an increasing sequence $\{\mathcal{G}_t\}$ of σ-algebras representing information flow. The σ-algebra $\mathcal{G}_t$, or *information set at time* t, contains events known up to time t. A *stopping time* τ is a random variable such that $\tau^{-1}(-\infty, t] \in \mathcal{G}_t$ for all t. A stopping time is therefore a random time that is known to the agent at time t. In our context, this means that the timing of interest rate adjustments that have occurred up to time t is known to the private sector at time t.

7. Evidence exists that inflation and output gap processes are persistent. In our model, this translates to the choice of modeling the output gap and inflation as continuous-time random walks. Relaxing the unit root restriction, one could allow for mean reversion. In continuous time, we can readily incorporate more general cost-push shocks by using an Ornstein–Uhlenbeck process, a sum of Brownian motion and a mean-reverting drift. Mean reversion in the cost-push shock would have two competing effects on the central bank's policy. The tendency of the process to revert to its long run mean makes the central bank more reluctant to adjust. On the other hand, mean reversion makes the process stationary, which in turn decreases central bank's uncertainty, thereby increasing the frequency of adjustment.

8. More formally, $\mathcal{A}$ acts on bounded Borel-measurable functions on $\mathbf{R}$.

9. The value function V, we construct below using smooth pasting, is actually not twice differentiable. The appendix contains a proof why this derivation is nevertheless true.

10. Dynkin's Formula is a classical result that generalizes the Fundamental Theorem of Calculus $V\big(z(\tau)\big) = V\big(z(0)\big) + \int_0^\tau V'\big(z(t)\big)z'(t)dt$ from ordinary differential equations to Itô diffusions [Protter (2004), p. 356].

11. The appendix shows the necessary first-order condition is in fact part of a sufficient condition, as assumption (iii) of the theorem.

12. Incorporating stochastic volatility into our model, i.e., allowing the central bank to anticipate possible changes in uncertainty regime, would retain the general adjustment/nonadjustment regions of optimal policy. The nonadjustment region would depend on the specification of the volatility stochastic process. Although a similar comparison between different specifications and their resulting policies can be carried out, the simplifying assumption of an unanticipatory central bank highlights the effects of uncertainty in our setting and unburdens the discussion from technicalities.

13. The hitting time distribution of Brownian motion can be derived using the strong Markov property. See, for example, Dixit (1993).

14. Since the problem is symmetric, all the intuitions that follow also applies to a deflation shock.

15. The early 80's corresponded to a period where the Fed adopted a money supply target instead of targeting the interest rate. Therefore, it is appropriate to set that period as the starting time of our analysis.

16. Galí and Gertler (1999) estimate the quarterly Philips curve

$$\pi_t = \kappa^q x_t + \beta^q \mathbf{E_t}(\pi_{\mathbf{t+1}})$$

with $\kappa^q = 0.023$, where superscript q denotes the quarterly parameter. Using the martingale assumption that $\pi_t = \mathbf{E_t}(\pi_{\mathbf{t+1}})$, we have

$$\pi_t = \frac{\kappa^q}{1-\beta^q} x_t = \frac{\kappa}{1-\beta} x_t.$$

This relationship between x_t and π_t holds regardless of data frequency. Thus, we can back out monthly κ by matching coefficients. With an annual discount rate of 4%, the implied monthly and quarterly discount rates are $\beta = 0.9967$, and $\beta^q = 0.99$, respectively. The implied monthly κ is 0.0076. For robustness check, when we also take the alternative value of 5.355% as the annual discount rate, the implied discount rates are $\beta^q = 0.9870$, and $\beta = 0.9956$, with $\kappa = 0.0078$.

17. Not only do we assume that volatility before and after 2001 is different, we also assume it is constant within each regime and not adapted to the central bank's filtration. Therefore, the central bank in our model cannot anticipate volatility changes ex ante. Incorporating stochastic volatility into the central bank's decision in our framework is a possible topic to be explored in detail in future research. More detailed remarks on this possibility are given in Section 8.

18. Since the New Keynesian Phillips curve is widely used to study monetary policy, and our purpose is not to test the Phillips curve, we condition on the validity of the model and identify cost-push shocks using the model. This back-solving strategy is also used in *The Conquest of American Inflation* [Sargent (1999)], where a series of residuals are inferred from the Phillips curve using observed data and the initial conditions for the government's beliefs.

19. For the sake of tractability, the volatility here is conservatively estimated by that of the cost-push shock, which is relatively smooth. One might expect that, in a richer model with more shocks, e.g., financial shocks, differences in volatility during recessions and normal regimes would be even larger, implying a stronger wait-and-see effect.

20. Lévy processes are continuous-time analogues of random walks, of which Brownian motion is a special case [Barndorff-Nielsen et al. (2001)]. One example is the independent sum of a Brownian motion and a compound Poission process. A general Lévy process can have infinitely many jumps during a finite time interval.

21. The *infinitesmal generator* L_z of an Itô diffusion z_t is the operator defined by, for $f : \mathbf{R^n} \to \mathbf{R}$ sufficiently smooth, $L_z f(z_0) = \lim_{\Delta t \to} \frac{E^{z_0}[f(z_{\Delta t})] - f(z_0)}{\Delta t}$. For example, L_z of the cost-push shock $dz_t = \theta \sigma dB_t$ is the second-order differential operator $\frac{1}{2}\theta^2\sigma^2\frac{d^2}{dz^2}$. For a central bank with value function V, $L_z V(z_0) = \frac{1}{2}\theta^2\sigma^2 V''(z_0)$ is therefore the expected marginal change of option value V under the law of dz_t with initial state z_0. This term, plus the flow cost, appears in the central bank's first-order condition (13).

22. See, for example, Øksendal (1985).

23. See Theorem 5.4 in Littman et al. (1963).

24. The notion of *distributional derivative* generalizes the classical derivative to *tempered distributions*, or *generalized functions* [see Gel'fand and Vilenkin (2014)].

25. Our proof strategy is as follows: First observe that V is C^2 except on a set of measure zero, where smooth pasting is applied. Its classical second derivative, defined almost everywhere, is in fact its distributional second derivative. Use the Sobolev Embedding Theorem to argue that the distributional second derivative can be well approximated by C^2-functions. Apply Dynkin's Formula to the approximating C^2-sequence and show that this is well behaved with respect to taking appropriate limit.

26. See, for example, Chapter 4 of Adams and Fournier (2003).

REFERENCES

Adams, Robert A. and John J. F. Fournier (2003) *Sobolev Spaces*. Amsterdam: Academic Press.

Alba, Joseph D. and Peiming Wang (2017) Taylor Rule and discretionary regimes in the United States: Evidence from a k-state Markov regime-switching model. *Macroeconomic Dynamics* 21, 817–833.

Alvarez, Fernando and Avinash Dixit (2014) A real options perspective on the future of the Euro. *Journal of Monetary Economics* 61, 78–109.

Barndorff-Nielsen, E. Ole, Thomas Mikosch, and Sidney I. Resnick (2001) *Lévy Processes: Theory and Applications*. New York: Springer.

Blinder, Alan S. (2009) Making monetary policy by committee. *International Finance* 12, 171–194.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry (2012) Really Uncertain Business Cycles. Technical report no. w18425, National Bureau of Economic Research.

Caplin, Andrew and John Leahy (1994) Business as usual, market crashes, and wisdom after the fact. *American Economic Review* 84, 548–565.

Davig, Troy and Eric M. Leeper (2008) Endogenous monetary policy regime change. In Lucrezia Reichlin and Kenneth West (eds.), *NBER International Seminar on Macroeconomics 2006*, pp. 345–391. Chicago: University of Chicago Press.

Dixit, Avinash K. (1993) *The Art of Smooth Pasting*. Switzerland: Harwood Academic Publishers.

Dixit, Avinash K. (1994) *Investment Under Uncertainty*. Princeton, NJ: Princeton University Press.

Dynkin, Evgeniĭ B. (1965) *Markov Processes*. New York: Springer.

Fernández-Villaverde, Jesús, Olaf Posch, and J. F. Rubio-Ramírez (2012) *Solving the new Keynesian model in continuous time*. Unpublished manuscript, Federal Reserve Bank of Atlanta.

Froot, Kenneth A. and Maurice Obstfeld (1989) Stochastic Process Switching: Some Simple Solutions. Technical report no. w2998, National Bureau of Economic Research.

Galí, Jordi (2009) *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton, NJ: Princeton University Press.

Galí, Jordi and Mark Gertler (1999) Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* 4, 195–222.

Gel'fand, Izrail M. and N. Ya. Vilenkin (2014) *Generalized Functions: Applications of Harmonic Analysis*. New York: Academic Press.

Holston, Kathryn, Thomas Laubach, and John, Williams (2017) Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*. doi:10.1016/j.jinteco.2017.01.004.

Kahn, George A. (2010) Taylor rule deviations and financial imbalances. *Federal Reserve Bank of Kansas City, Economic Review* 95(2), 63–99.

Levin, Andrew T. and John C. Williams (2003) Robust monetary policy with competing reference models. *Journal of Monetary Economics* 50, 945–975.

Levin, Andrew T., Volker Wieland, and John C. Williams (1999) Robustness of simple monetary policy rules under model uncertainty. In John B. Taylor (ed.), *Monetary Policy Rules*, pp. 263–318. Chicago: University of Chicago Press.

Lipster, Robert S. and Albert N. Shiryayev (1989) *Theory of Martingales*. Pordrecht, The Netherlands: Kluwer.

Littman, Walter, Guido Stampacchia, and Hans F. Weinberger (1963) Regular points for elliptic equations with discontinuous coefficients. *Annali della Scuola Normale Superiore di Pisa-Classe di Scienze* 17, 43–77.

Mankiw, N. Gregory and Ricardo Reis (2002) Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics* 117, 1295–1328.

Miao, Jianjun and Neng Wang (2011) Risk, uncertainty, and option exercise. *Journal of Economic Dynamics and Control* 35, 442–461.

Murray, Christian J., Alex Nikolsko-Rzhevskyy, and David H. Papell (2015) Markov switching and the Taylor principle. *Macroeconomic Dynamics* 19, 913–930.

Øksendal, Bernt (1985) Stochastic processes, infinitesimal generators and function theory. *Operators and Function Theory* 153, 139–162.

Protter, Philip E. (2004) *Stochastic Integration and Differential Equations*. Berlin: Springer.

Rotemberg, Julio J. and Michael Woodford (1999) Interest rate rules in an estimated sticky price model. In John B. Taylor (ed.), *Monetary Policy Rules*, pp. 57–126. Chicago: University of Chicago Press.

Rudebusch, Glenn D. (2006) Monetary policy inertia: Fact or fiction? *International Journal of Central Banking* 2, 85–135.

Sargent, Thomas J. (1999) *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.

Sims, Christopher A. (2010) Rational inattention and monetary economics. In Benjamin M. Friedman and Michael Woodford (eds.), *Handbook of Monetary Economics*, 1st ed., vol. 3, pp. 155–181. Amsterdam: North Holland.

Stokey, Nancy L. (2008) *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton, NJ: Princeton University Press.

Stokey, Nancy L. (2016) Wait-and-see: Investment options under policy uncertainty. *Review of Economic Dynamics* 21, 246–265.

Svensson, Lars E. O. and Noah Williams (2008) Optimal monetary policy under uncertainty: A Markov jump-linear-quadratic approach. *Federal Reserve Bank of St. Louis Review* 90, 275–293.

Taylor, John B. (2014) The role of policy in the Great Recession and the weak recovery. *American Economic Review* 104, 61–66.

Trojanowska, Magdalena and Peter M. Kort (2010) The worst case for real options. *Journal of Optimization Theory and Applications* 146, 709–734.

Woodford, Michael (1999) Optimal monetary policy inertia. *The Manchester School* 67, 1–35.

Woodford, Michael (2003) *Interest and Prices: Foundations of A Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

APPENDIX

This section contains a proof that the derivation contained in Section 4 holds when the value function V only satisfies the Euler equations, or so-called smooth pasting condition, rather than being C^2, i.e., twice differentiable with continuous second derivative. The difficulty is that the classical Itô's Lemma, and its corollary Dynkin's Formula, is applicable only for C^2-functions. A value function V constructed via smooth pasting is clearly not C^2 everywhere, in particular at points where the pasting is done. On its state space **R**, the central bank's value function V is only C^1; differentiating twice, as was done in (12) in Section 4, might not make sense. Here, we show that our results remain true nevertheless. In addition to the current model, this formal argument is provided also with a view toward extending our analysis to more general settings discussed in the Conclusion, such as one where the central bank is allowed to switch between periods of continuous adjustment and no adjustment, or is restricted by a zero lower bound, or faces a cost-push shock that is a more general Lévy process with jumps, or a stochastic volatility model. In practice, our theorem is a general verification theorem, i.e., one that provides a sufficient condition—the Euler equation—for optimality.

We recall the following notation: For $\mathbf{R^n}$-valued Itô diffusion z_t with infinitesimal generator L_z and a bounded domain $W \subset \mathbf{R^n}$, consider the expectations operator $f(z) \mapsto E^z[\int_0^{\tau_V} L_z f(z_t) dt]$[21]. If $E^z[\tau_V] < \infty$ for all $z \in W$, then a finite measure, called the *Green measure*, can be defined on W by[22]

$$\int_V \phi(z')d\mu_z(z') = E^z\left[\int_0^{\tau_V} \phi(z_t)dt\right], \quad \forall \phi \in C(\overline{W}).$$

In other words, the Green measure μ_z is the length of time that z_t, given the initial state z, is expected to spend in W. In our model, the W of interest is the conjectured inaction region $(a, b) \subset \mathbf{R}$ of the central bank.

If the differential operator L_z is uniformly elliptic on D, then for all $z \in D$, μ_z is absolutely continuous with respect to Lebesgue measure, in which case we denote its density by $d\mu_z(x') = G(z, z')dz'$. In fact, one has[23]

$$G(z, z') \in L^q_{\text{loc}}(\mathbf{R^n}), \quad \text{if } \mathbf{q} < \mathbf{1} + \frac{\mathbf{n}}{\mathbf{2}},$$

where $L^q_{\text{loc}}(\mathbf{R^n})$ denotes the Lebesgue space of functions with finite L^q-norm when restricted to compact subsets.

As stated in Section 4, the central bank's problem can be reduced to an optimal stopping problem. Therefore, we prove sufficiency of the Euler equation in terms of the general optimal stopping problem. Let $g : \mathbf{R^n} \to \mathbf{R}$ be Borel-measurable. The problem is to find

$$V(z) = \inf_\tau E^z[g(z_\tau)],$$

where $\inf_\tau$ denotes the infimum over all stopping times. In our context, g is the central bank's time discounted value function plus total flow cost due to deviations driven by the cost-push shock.

THEOREM. *Let $V : \mathbf{R^n} \to \mathbf{R}$ that satisfy the following conditions:*

(i) There exists a region $D \subset \mathbf{R^n}$ with C^1-boundary δD and $V \in C^2_b(\mathbf{R^n} \backslash \delta D)$.
(ii) $h|_D \geq g|_D$. (Continuation)
(iii) $L_z h = 0$ on D, where L_z is the infinitesimal generator of z_t. (First-order condition)
(iv) $h|_{\mathbf{R^n}-D} = g|_{\mathbf{R^n}-D}$ and $L_X h \leq 0$ on $\mathbf{R^n} - \overline{D}$. (Noncontinuation)
(v) $h \in C^1(\mathbf{R^n})$. (Euler equation/C^1-smooth pasting condition)

Then, V solves the optimal stopping problem, with the infimum being attained by the first exit time of D.

Proof. Using Assumptions (i) (C^1-boundary and C^2 almost everywhere), and (v) (C^1-pasting), integration by parts shows that the second mixed partials $\delta_{z_i z_j} h$, defined almost everywhere (outside δD), is the distributional second derivative of V.[24,25]

By the boundedness assumption on V [part of (i)] and existence of distribution derivatives, V lies in the L^p-Sobolev space $W^{2,p}$ for any p. Since $C^2(\mathbf{R^n})$ is dense in $W^{2,p}$, there exists a sequence $\{f_k\} \subset C^2(\mathbf{R^n})$, such that

$$\|f_k - h\|_{W^{2,p}} \to 0.$$

On the other hand, by the Sobolev Embedding Theorem, if $p > \frac{n}{2}$,[26]

$$\|f_k - h\|_\infty \to 0.$$

Now, Dynkin's Formula holds for each f_k, that is,

$$E^z[f_k(z_\tau)] = f_k(z) + \int_V L_z f_k(z')G(z, z')dz'.$$

So,

$$\lim_{k\to\infty} E^z[f_k(z_\tau)] = \lim_{k\to\infty} f_k(z) + \lim_{k\to\infty} \int_V L_z f_k(z')G(z,z')dz'.$$

Choose $p > \frac{n}{2} + 1$, then its conjugate exponent q satisfies $q < 1 + \frac{2}{n}$. The boundedness assumption means that the Sobolev norm can be approximated by the uniform norm:

$$\lim_{k\to\infty} f_k(z) = V(z).$$

As a by-product, the approximating sequence $\{f_k\}$ can be chosen as regular, in the sense of large p, to accommodate possible singularities of the Green density: by choice of $\{f_k\}$,

$$\|L_z f_k(z') - L_z V(z')\|_{L^p} \to 0.$$

By Hölder's inequality,

$$|\int_V L_z f_k(z') - L_z V(z')G(z,z')dz'| \leq \|L_z f_k(z') - L_z V(z')\|_{L^p} \cdot \|G(z,z')\|_{L^q} \to 0.$$

(Strictly speaking, we have assumed D is compact, but this is without loss of generality). So, for all z,

$$\lim_{k\to\infty} E^z[f_k(z_\tau)] = V(z) + \int_V L_z V(z')G(z,z')dz'.$$

By the Dominated Convergence Theorem,

$$E^z[V(z_\tau)] = \lim_{k\to\infty} E^z[f_k(z_\tau)].$$

So, we have that a Dynkin's Formula holds for V:

$$E^z[V(z_\tau)] = V(z) + \int_W L_z V(z')G(z,z')dz'.$$

It follows that h is superharmonic locally on δD, therefore everywhere [see, Dynkin (1965), p. 22]. This proves the theorem. ■