Gradualism in Monetary Policy: A Time-Consistency Problem?*

Jeremy C. Stein
Harvard University and NBER

Adi Sunderam
Harvard University and NBER

June 2015

Abstract

We develop a model of monetary policy with two key features: (i) the central bank has private information about its long-run target for the policy rate; and (ii) the central bank is averse to bond-market volatility. In this setting, discretionary monetary policy is gradualist, or inertial, in the sense that the central bank only adjusts the policy rate slowly in response to changes in its privately-observed target. Such gradualism reflects an attempt to not spook the bond market. However, this effort ends up being thwarted in equilibrium, as long-term rates rationally react more to a given move in short rates when the central bank moves more gradually. The same desire to mitigate bond-market volatility can lead the central bank to lower short rates sharply when publicly-observed term premiums rise. In both cases, there is a time-consistency problem, and society would be better off appointing a central banker who cares less about the bond market. We also discuss the implications of our model for forward guidance once the economy is away from the zero lower bound.

* We thank Adam Wang-Levine for research assistance, and seminar participants at Stanford for their comments.

I. Introduction

In this paper, we attempt to shed some light on the well-known phenomenon of gradualism in monetary policy. As described by Bernanke (2004), gradualism is the idea that "the Federal Reserve tends to adjust interest rates incrementally, in a series of small or moderate steps in the same direction." This behavior can be represented empirically by an inertial Taylor rule, with the current level of the federal funds rate modeled as a weighted average of a target rate—which is itself a function of either contemporaneous or expected values of inflation and the output gap as in, e.g., Taylor (1993)—and the lagged value of the funds rate. In this specification, the coefficient on the lagged funds rate captures the degree of inertia in policy. In recent U.S. samples, estimates of the degree of inertia are strikingly high, on the order of 0.85 in quarterly data.[1]

Several authors have proposed theories in which this kind of gradualism can be thought of as optimal behavior on the part of the central bank. One influential line of thinking, due originally to Brainard (1967) and refined by Sack (1998), is that moving gradually makes sense when there is uncertainty about how the economy will respond to a change in the stance of policy. An alternative rationale for gradualism comes from Woodford (2003), who argues that if the central bank can commit itself to following an inertial rule, doing so gives it more leverage over long-term rates for a given change in short rates, a property which is desirable in the context of his model.

In what follows, we offer a different take on gradualism. In our model, the observed degree of policy inertia is not optimal from an ex ante perspective, but rather reflects a fundamental time-consistency problem. The distinguishing feature of our approach is that the Fed is assumed to have private information about its preferred value of the target rate. In other words, the Fed knows something about its reaction function that the market does not. We also assume that the Fed behaves as if it is averse to bond-market volatility. We model this concern in reduced form, by simply putting the volatility of long-term rates into the Fed's objective function, but stress that a preference of this sort can ultimately be rooted in an effort to deliver on its traditional dual mandate. For example, a

[1] Coibion and Gorodnichenko (2012) provide a comprehensive recent empirical treatment; see also Rudebusch (2002, 2006). Campbell, Pflueger and Viceira (2015) argue that the degree of inertia in Federal Reserve rate-setting became more pronounced after about 2000.

bout of bond-market volatility may be undesirable not just in its own right, but rather because it is damaging to the financial system and hence to real economic activity and employment.

Nevertheless, in a world of private information and discretionary meeting-by-meeting decision-making, an attempt by the Fed to moderate bond-market volatility can be welfare-reducing. The logic here is similar to that in signal-jamming models (Holmstrom, 1999; Stein, 1989). Suppose the Fed observes a private signal that its long-run target value for the funds rate has permanently increased by 100 basis points. If it adjusts policy fully in response to this signal, raising the funds rate by 100 basis points, long-term rates will move by a similar amount. If it is averse to such a movement in long-term rates, the Fed will be tempted to announce a smaller change in the funds rate, thereby trying to fool the market into thinking that its private signal was not as dramatic. Hence it will under-adjust to its signal, perhaps raising the funds rate by only 25 basis points.

However, if bond-market investors come to understand this dynamic, the Fed's efforts to reduce volatility will be frustrated in equilibrium. The market will see the 25 basis-point increase in the funds rate and understand that it is likely to be just the first in a series of similar moves, so long-term rates will react more than one-for-one to the change in short rates. Indeed, in a rational-expectations equilibrium, the Fed's private signal will always be fully inferred by the market, regardless of the degree of gradualism. Still, if it acts on a discretionary basis, the Fed will always keep trying to fool the market. This is because when it decides how much to adjust the policy rate, it *takes as given the market's conjecture* about the degree of inertia in its rate-setting behavior. As a result, the Fed's behavior is inefficient from an ex ante perspective: by moving gradually, it does not succeed in its attempts to reduce bond-market volatility, but gradualism means that the policy rate is further from its long-run target than it otherwise would be.

This inefficiency reflects a commitment problem. In particular, the Fed cannot commit to not trying to smooth the private information that it communicates to the market via its changes in the policy rate.[2] One institutional solution to this problem, in the spirit of Rogoff (1985), would be to appoint a central banker who cares less about bond-market volatility than the representative member

---

[2] The literature on monetary policy has long recognized a different commitment problem, namely that, under discretion, the central bank will be tempted to create surprise inflation so as to lower the unemployment rate. See, e.g., Kydland and Prescott (1977) and Barro and Gordon (1983). More recently, Farhi and Tirole (2012) have pointed to the time-consistency problem that arises from the central bank's ex post desire to ease monetary policy when the financial sector is in distress; their focus on the central bank's concern with financial stability is somewhat closer in spirit to ours.

of society. More broadly, appointing such a market-insensitive central banker can be thought of as a metaphor for building a certain kind of institutional culture and set of norms inside the central bank such that bond-market movements are not given as much weight in policy deliberations.

We begin with a simple static model that is designed to capture the above intuition in as parsimonious a way as possible. The main result here is that in any rational-expectations equilibrium, there is always under-adjustment of the policy rate, as compared to a first-best outcome in which the Fed adjusts the policy rate fully in response to changes in its privately-observed target. Moreover, for some parameter values, there can be Pareto-ranked multiple equilibria with different degrees of under-adjustment. The intuition for these multiple equilibria is that there is a two-way feedback between the market's expectations about the degree of gradualism on the one hand and the Fed's optimal choice of gradualism on the other.[3] Specifically, if the market conjectures that the Fed is behaving in a highly inertial fashion, it will react more strongly to an observed change in the policy rate: in an inertial world, the market knows that there are further changes to come. But this strong sensitivity of long-term rates to changes in the policy rate makes the Fed all the more reluctant to move the policy rate, hence validating the initial conjecture of extreme inertia.

Next, we ask whether the specific inefficiency that we have identified can be mitigated with forward guidance, whereby the Fed announces in advance a preferred path for future short rates. We show that it cannot. Because its future private information is by definition not forecastable, forward guidance is a one-sided commitment device. It can only commit the Fed to incorporating new private information more slowly, not more quickly. Moreover, we argue that once the economy is away from the zero lower bound, forward guidance can actually be harmful if it is not implemented carefully. If we are in a region of the parameter space where there are multiple equilibria, forward guidance can increase the risk of getting stuck in the Pareto-inferior more-gradualist equilibrium.

We then enrich the model by adding term premium shocks as a source of publicly observable variation in long-term interest rates. We show that in its efforts to reduce bond-market volatility, the Fed is tempted to try to offset term premium shocks, cutting the policy rate when term premiums rise. Effectively, the Fed tries to pretend that it has received dovish private information about the optimal target rate whenever there is a positive term premium shock. Again, it does this even though in

---

[3] For an informal description of this two-way feedback dynamic, see Stein (2014).

equilibrium, this tactic is unsuccessful in reducing volatility. Thus, the presence of term premium shocks exacerbates the Fed's time-consistency problem.

Finally, we extend the model to an explicitly dynamic setting, which allows us to more fully characterize the impact of the Fed's gradualist behavior on the entire term structure of interest rates. We also outline the conditions under which the key qualitative conclusions from our static model carry over to the dynamic version.

The remainder of the paper is organized as follows. Section II discusses some motivating evidence, based on readings of Federal Open Market Committee (FOMC) transcripts. Section III presents the static version of the model and summarizes our basic results on under-adjustment of the policy rate, multiple equilibria, forward guidance, and term premiums. Section IV develops a dynamic extension of the model, and Section V concludes.

II. Motivating Evidence from FOMC Transcripts

In their study of monetary-policy inertia, Coibion and Gorodnichenko (2012) use FOMC transcripts to document two key points. First, FOMC members sometimes speak in a way that is suggestive of a gradual-adjustment model—that is, they articulate a target for the policy rate, and then put forward reasons why it is desirable to adjust only slowly in the direction of that target. Second, one of the stated rationales for such gradualism appears to be a desire not to create financial-market instability. Coibion and Gorodnichenko highlight the following quote from Chairman Alan Greenspan at the March 1994 FOMC meeting:

> "My own view is that eventually we have to be at 4 to 4½ percent. The question is not whether but when. If we are to move 50 basis points, I think we would create far more instability than we realize, largely because a half-point is not enough to remove the question of where we are ultimately going. I think there is a certain advantage in doing 25 basis points…."

In a similar spirit, at the August 2004 meeting, shortly after the Fed had begun to raise the funds rate from the low value of one percent that had prevailed since mid-2003, Chairman Greenspan remarked:

> "Consequently, the sooner we can get back to neutral, the better positioned we will be. We were required to ease very aggressively to offset the events of 2000 and 2001, and we took the funds rate down to extraordinarily low levels with the thought in the back of our minds, and often said around this table, that we could always reverse our actions. Well, reversing is not all that easy….We've often discussed that ideally we'd like to be in a position where, when we move as we did on June 30 and I hope today, the markets respond with a shrug. What that means is that the

adjustment process is gradual and does not create discontinuous problems with respect to balance sheets and asset values."

These sorts of quotes help motivate our basic modeling approach, in which gradualism in monetary policy reflects the Fed's desire to keep bond-market volatility in check—in Greenspan's words, to "not create discontinuous problems with respect to balance sheets and asset values." This same approach may also be helpful in thinking about changes in gradualism over time. Campbell, Pflueger and Viceira (2015) show that the degree of inertia in Fed rate-setting behavior became significantly more pronounced after about 2000; given the logic of our model, one might wonder whether this heightened inertia was associated with an increase over time in the Fed's concern with financial markets. In a crude attempt to speak to this question, we examine all 200 FOMC transcripts for the 25-year period 1985-2009, and for each meeting, simply measure the frequency of words related to financial markets.[4] Specifically, we count the number of times the terms "financial market", "equity market", "bond market", "credit market", "fixed income", and "volatility" are mentioned.[5] For each year, we aggregate this count, and divided by the total number of words in that year's transcripts.

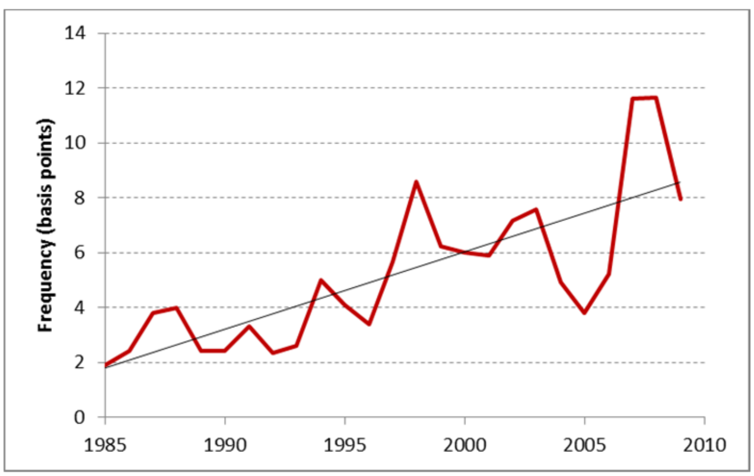

Figure 1. Frequency of terms related to financial markets in FOMC transcripts.

[4] Given the five-year lag in making transcripts public, 2009 is the last available year.

[5] We obtain similar results if use different subsets of these terms. For instance, the results are similar if we only count the frequency of the term "financial market".

Figure 1 displays the results from this exercise. As can be seen, there is a strong upward trend in the data. While there is also a good deal of volatility overlaid on top of the trend— mentions of financial-market terms decline in the years preceding the financial crisis and then spike up—a simple linear time trend captures almost 59 percent of the variation in the time series, with a t-statistic of 5.0.[6] And the fitted value of the series based on this time trend goes up about four-fold over the 25-year sample period. While this analysis is extremely simplistic, it does suggest an increasing emphasis over time by the FOMC on financial-market considerations. If there was in fact such a change in FOMC thinking, the model we develop below is well-suited to drawing out its implications, both for the dynamics of the policy rate, as well as for longer-term yields.

III. Static Model

We begin by presenting what is effectively a static version of the model, in which the Fed adjusts the funds rate only partially in response to a one-time innovation in its desired target rate. In this static version, the phenomenon of "gradualism" is therefore really nothing more than partial adjustment. Later, in Section IV, we extend the model to a fully dynamic setting, in which we can be more explicit about how an innovation to the Fed's target rate at any time t gradually makes its way into the funds rate over the following sequence of dates. In spite of its limitations on the realism dimension, the static model highlights the main intuition for why a desire to limit bond-market volatility creates a time-consistency problem for the Fed.

A. Model Setup

We begin by assuming that at any time t, the Fed has a target rate based on its traditional dual-mandate objectives. That is, this target rate is the Fed's best estimate of the value of the federal funds rate that keeps inflation and unemployment as close as possible to their desired levels. For tractability, we assume that the log of the target rate, denoted by i_t^*, follows a random walk, so that:

$$i_t^* = i_{t-1}^* + \varepsilon_t, \tag{1}$$

---

[6] Restricting the sample to the pre-crisis period 1985-2006, we obtain an R^2 of 46 percent and a t-statistic of 4.3. So the trend in the data is not primarily driven by post-2006 period.

where $\varepsilon_t \sim N\left(0, \sigma_\varepsilon^2 \equiv \frac{1}{\tau_\varepsilon}\right)$ is normally distributed. Our key assumption is that i_t^* is private information of the Fed and is unknown to market participants before the Fed acts at time t. One can think of the private information embodied in i_t^* as arising from the Fed's attempts to follow something akin to a Taylor rule, but where it has private information about either the appropriate coefficients to use in the rule (i.e., its reaction function), or about its own forecasts of future inflation or unemployment. At some level, an assumption of private information along these lines is necessary if one wants to understand why asset prices respond to Fed policy announcements.[7]

Once it knows the value of i_t^*, the Fed acts to incorporate some of its new private information ε_t into the federal funds rate i_t, which is observable to the market. We assume that the Fed picks i_t on a discretionary period-by-period basis to minimize the loss function L_t, given by:

$$L_t = \left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^\infty\right)^2, \tag{2}$$

where i_t^∞ is the log of the infinite-horizon forward rate. Thus, the Fed has the usual concerns about inflation and unemployment, as captured in reduced form by a desire to keep i_t^* close to i_t. However, when $\theta > 0$, the Fed also cares about the volatility of long-term bond yields, as measured by the squared change in the infinite-horizon forward rate.

For simplicity, we start by assuming that the log expectations hypothesis holds, so there is no time-variation in the term premium. Then, because the log of the target rate i_t^* follows a random walk, the log of the distant forward rate is given by the market's expectation of i_t^* at time t. Specifically, we have $i_t^\infty = E_t\left[i_t^*\right]$. In Section III.E below, we relax this assumption so that we can also consider the Fed's reaction to term-premium shocks.

Several features of the Fed's loss function are worth discussing. First, in our simple formulation, θ reflects the degree to which the Fed cares about bond-market volatility, over and above its desire to keep i_t^* close to i_t. However, this loss function should not be taken to mean that

---

[7] Said differently, if the Fed mechanically followed a policy rule that was a function only of publicly observable variables (e.g., the current values of the inflation rate and the unemployment rate), then the market would react to news releases about movements in these variables but not to FOMC policy statements.

the Fed behaves as if it has a triple mandate which includes a separate financial-stability objective—i.e., we do not intend to suggest that the Fed cares about asset prices for their own sake. Rather, the loss function is meant to capture the idea that volatility in financial market conditions can affect the real economy and hence the Fed's ability to satisfy its traditional dual mandate. This linkage is not modeled explicitly here, but as one example of what we have in mind, the Fed might be concerned that a spike in bond-market volatility could damage highly-levered intermediaries and interfere with the credit-supply process. With this stipulation in mind, we take as given that θ reflects the socially "correct" objective function—in other words, it is exactly the value that a well-intentioned social planner would choose. We then ask whether there is a time-consistency problem when the Fed tries to optimize this objective function period-by-period, in the absence of any kind of commitment technology.

Second, it might seem more natural to think of the Fed as caring about the volatility of bond prices rather than forward rates. But since we are working with logs, caring about rate volatility is actually similar to caring about price volatility since $p_t^n = -ny_t^n$, where p_t^n is the log price of an n-period bond at time t, and y_t^n is its log yield.[8] In a related vein, one could imagine that the Fed cares about the *level* of the long-term bond rate at time t rather than its volatility. This is a meaningful difference in the sort of dynamic model that we explore in Section IV. However, in our simple static setup, there is no distinction between the two, as there is just a single shock that both pushes the level of the rate away from its steady state value and simultaneously induces volatility in the rate.

Finally, the assumption that the Fed cares about the infinite-horizon forward rate i_t^∞, along with the assumption that the target rate i_t^* follows a random walk, greatly simplifies both the Fed's and the market's problem. Because i_t^* follows a random walk, all of the new private information ε_t will, in expectation, eventually be incorporated into the future short rate. The reaction of the infinite-horizon forward rate reflects this revision in expectations and allows us to abstract from the exact dynamic path that the Fed follows in ultimately incorporating its new information into the funds rate. In contrast, revisions in finite-horizon forward rates depend on the exact path of the short rate. For

[8] The correspondence is, however, not exact since we are assuming the Fed cares about forward rates, not bond yields.

example, suppose i_t^* were public information and changed from 1% to 5%. The infinite-horizon forward rate would then immediately jump to 5%, while the two-year forward rate might move by less if the market expected the Fed to take a long time to implement the tightening. We relax the assumption that the Fed cares about the infinite-horizon forward rate and also work out the full adjustment path of rates to an innovation in i_t^*, when we present the dynamic model in Section IV.

Once it observes i_t^*, we assume that the Fed sets the federal funds rate i_t by following a partial adjustment rule of the form:

$$i_t = i_{t-1} + k\left(i_t^* - i_{t-1}\right) + u_t, \tag{3}$$

where $u_t \sim N\left(0, \sigma_u^2 \equiv \frac{1}{\tau_u}\right)$ is normally distributed noise that is overlaid onto the rate-setting process. The noise u_t is nothing more than a technical modeling device; its usefulness will become clear shortly. Loosely speaking, this noise—which can be thought of as a tiny "tremble" in the Fed's otherwise optimally-chosen value of i_t —ensures that the Fed's actions cannot be perfectly inverted to fully recover its private information i_t^*. As will be seen, this imperfect-inversion feature helps to avoid some degenerate equilibrium outcomes. In any event, in all the cases we consider, σ_u^2 is set to be very small relative to σ_ε^2, often several orders of magnitude smaller, so that one should not try to read too much economic content into the construct. One possibility would be to think of u_t as coming from the discreteness associated with the fact that the Fed uses round numbers (typically in 25 basis-point increments) for the funds rate settings that it communicates to the market, while its underlying private information about i_t^* is presumably continuous.

The market seeks to infer the Fed's private information i_t^* based on its observation of the funds rate i_t. To do so, it conjectures that the Fed follows a rule given by:

$$i_t = i_{t-1} + \kappa\left(i_t^* - i_{t-1}\right) + u_t. \tag{4}$$

Thus, the market correctly conjectures the form of the Fed's smoothing rule but, crucially, it does not directly observe the Fed's smoothing parameter *k*; rather, it has to make a guess κ as to the value of this parameter. In a rational-expectations equilibrium, this guess will turn out to be correct, and we will have $k = \kappa$. However—and this is a key piece of intuition—when the Fed decides how much smoothing to do, it takes the market's conjecture κ as a fixed parameter and does not impose that $k=\kappa$. The equilibrium concept here is thus of the "signal-jamming" type introduced by Holmstrom (1999): the Fed, taking the bond market's estimate of κ as fixed, tries to fool the market into thinking that i_t^* has moved by less than it actually has, in an effort to reduce the volatility of long-term rates. In a Nash equilibrium, these attempts to fool the market wind up being fruitless, but the Fed can't resist the temptation to try.[9]

B. Equilibrium

We are now ready to solve for the equilibrium in the static model. Suppose that the economy was previously in steady state at time *t*-1, with $i_{t-1}=i_{t-1}^*$. Given the Fed's adjustment rule, the funds rate at time *t* satisfies:

$$i_t = i_{t-1} + k\left(i_t^* - i_{t-1}\right) + u_t = i_{t-1} + k\varepsilon_t + u_t. \tag{5}$$

Based on its conjecture about the Fed's adjustment rule in equation (4), the market tries to back out i_t^* from its observation of i_t. Since both shocks ε_t and u_t are normally distributed, the market's expectation is given by:

$$E\left[i_t^* \mid i_t\right] = i_{t-1} + \frac{\kappa\tau_u \times \left(i_t - i_{t-1}\right)}{\tau_\varepsilon + \kappa^2\tau_u}. \tag{6}$$

The less noise there is in the Fed's adjustment rule, the higher is τ_u and the more the market reacts to the change in the rate i_t.

In light of equations (5) and (6) and the random-walk property that $i_t^\infty = E_t\left[i_t^*\right]$, the Fed's loss function taking expectations over the realization of the noise can be written as:

$$L_t = E_{u_t}\left[\left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^\infty\right)^2\right] = \left(1-k\right)^2\varepsilon_t^2 + \sigma_u^2 + \theta\chi^2\left(k^2\varepsilon_t^2 + \sigma_u^2\right), \tag{7}$$

[9] There is a close analogy to models where corporate managers with private information pump up their reported earnings in an effort to impress the stock market. See, e.g., Stein (1989).

where $\chi = \dfrac{\kappa\tau_u}{\tau_\varepsilon + \kappa^2\tau_u}$. The Fed then minimizes this loss function by picking the optimal value of k, where, once again, we emphasize that in doing so, it takes the market's conjectures about its behavior, and hence the parameter χ, as fixed. The first-order condition with respect to k then yields:

$$k = \frac{1}{1+\theta\chi^2}. \tag{8}$$

In rational-expectations equilibrium, the market's conjecture turns out to be correct, so we have $\kappa = k$. Imposing this condition, we have that in equilibrium, the Fed's adjustment rule satisfies:

$$k = \frac{\left(\tau_\varepsilon + k^2\tau_u\right)^2}{\left(\tau_\varepsilon + k^2\tau_u\right)^2 + \theta\left(k\tau_u\right)^2}. \tag{9}$$

The following proposition characterizes the partial-adjustment nature of any rational expectations equilibrium.

Proposition 1: In any rational expectations equilibrium, the Fed's adjustment to a change in its target rate is partial: we have $k < 1$ so long as $\theta > 0$.

Proof: All proofs are given in the Appendix.

B.1 Equilibrium with No Noise in Rate-Setting Process

To build some intuition for the Fed's behavior, let us begin by considering the simple limiting case in which there is no noise in the rate setting process: $\sigma_u^2 = 1/\tau_u = 0$. In this case, the market's inference of i_t^* is simply:

$$E\left[i_t^* \mid i_t\right] = i_{t-1} + \frac{\left(i_t - i_{t-1}\right)}{\kappa} = i_{t-1} + \frac{k}{\kappa}\varepsilon_t, \tag{10}$$

and the Fed's loss function is:

$$L_t = \left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^\infty\right)^2 = \left(\left(1-k\right)\varepsilon_t\right)^2 + \theta\left(\frac{k}{\kappa}\varepsilon_t\right)^2. \tag{11}$$

When the Fed considers lowering k, it trades off being further away from the optimal i_t^* against the fact that it believes it can reduce bond market volatility by moving more slowly. In rational-expectations equilibrium, equation (9) reduces to:

$$k = \frac{k^2}{k^2 + \theta}. \tag{12}$$

First note that when $\theta = 0$, the only solution to (12) is given by $k = 1$. When the Fed does not care about bond market volatility, it fully adjusts to its new private information ε_t. By contrast, when $\theta > 0$, there may be more than one solution to (12), but in any equilibrium it must be the case that $k < 1$, and the Fed only partially adjusts.

Moreover, when $\theta > 0$, it is always the case that $k = 0$ satisfies equation (12), and there is therefore one equilibrium in which *the Fed does not adjust the funds rate at all*. This somewhat unnatural outcome is a function of the extreme feedback that arises between the Fed's adjustment rule and the market's conjecture when there is no noise in the rate-setting process. Specifically, when the market conjectures that the Fed never moves the funds rate at all (i.e., when the market conjectures that $\kappa = 0$), then even a tiny out-of-equilibrium move by the Fed would lead to an infinitely large change in the infinite-horizon forward rate. Anticipating this possibility, the Fed validates the market's conjecture by not touching the dial at all, i.e. by choosing $k = 0$. However, as soon as there is even an infinitesimal amount of noise u_t in the rate-setting process, this extreme $k = 0$ equilibrium is ruled out, as small changes in the funds rate now lead to bounded market reactions. This explains our motivation for keeping a tiny amount of rate-setting noise in the more general model.

In addition to $k = 0$, for $0 < \theta < 0.25$ there are two other solutions to equation (12), given by:

$$k = \frac{1 \pm \sqrt{(1 - 4\theta)}}{2} \tag{13}$$

Of these, the larger of the two values also represents a stable equilibrium outcome. Thus, as long as θ is not too big, the model also admits a non-degenerate equilibrium, with ½ < k < 1, even with $\sigma_u^2 = 0$. And within this region, higher values of θ lead to lower values of k. In other words, the more intense the Fed's concern with bond-market volatility, the more gradually it adjusts the funds rate.

B.2 Equilibrium Outcomes across the Parameter Space

Now we return to the more general case where $\sigma_u^2 > 0$ and explore more fully the range of outcomes produced by the model for different parameter values. In each panel of Figure 2 below, we plot the Fed's best response k as a function of the market's conjecture κ. Any rational expectations equilibrium must lie on the 45-degree line where $k = \kappa$.

In Panel (a) of the figure, we begin with a relatively low value of θ, namely 0.12, and set $\sigma_\varepsilon^2 / \sigma_u^2 = 250$. As can be seen, this leads to a unique equilibrium with a high value of k, given by 0.86. When the Fed cares only a little bit about bond-market volatility, it engages in a modest amount of rate smoothing, because it does not want to deviate too far from its target rate i_t^* in its efforts to reduce volatility.

In Panel (b), we keep $\sigma_\varepsilon^2 / \sigma_u^2 = 250$ and increase θ to 0.40. In this high-θ case, there is again a unique equilibrium, but it now involves a very low value of k of just 0.04. Thus, when the Fed cares a lot about bond market volatility, the funds rate dramatically under-adjusts to its new information.

In Panel (c), we set θ to an intermediate value of 0.22. Here, we have multiple equilibria: the Fed's best response crosses the 45-degree line in three places. Of these three crossing points, the two outer ones (at k = 0.07 and 0.69 respectively) correspond to stable equilibria, in the sense that if the market's initial conjecture κ takes on an out-of-equilibrium value, the Fed's best response to that conjecture will tend to drive the outcome towards one of these two extreme crossing points.

The existence of multiple equilibria highlights an essential feature of the model: the potential for market beliefs about Fed behavior to become self-validating. If the market conjectures that the Fed will adjust rates only very gradually, then even small changes are heavily freighted with informational content about the Fed's reaction function. Given this strong market sensitivity and its desire not to create too much volatility, the Fed may then choose to move very carefully, even at the cost of accepting a funds rate that is quite distant from its target. Conversely, if the market conjectures that the Fed is more aggressive in its adjustment of rates, it reacts less to any given movement, which frees the Fed up to track its target rate more closely.

When there are multiple equilibria, it is typically the case that the one with the higher value of k leads to a better outcome from the perspective of the Fed's objective function: with a higher k, i_t is closer to i_t^*, and yet in equilibrium, bond-market volatility is no greater. This suggests that if we are in a range of the parameter space where multiple equilibria are a possibility, it is important for the Fed to avoid doing anything in its communications that tends to lead the market to expect an overly low value of k. Even keeping its own preferences fixed, fostering an impression of strong gradualism among market participants can lead to an undesirable outcome in which the Fed gets stuck in the low-k equilibrium. We will return to this point in more detail shortly, when we discuss the implications of forward guidance in our model.

As noted earlier, the very small but non-zero value of σ_u^2 plays an important role in pinning down the "low-k" equilibria illustrated in Panels (b) and (c)—those where we have $k < 0.5$. If instead we were to set $\sigma_u^2 = 0$, these equilibria would degenerate to $k = 0$, for the reasons described above. By contrast, the "high-k" equilibrium depicted in Panel (a) as well as the upper equilibrium in Panel (c) are much less sensitive to the choice of σ_u^2. Thus, loosely speaking, the value of σ_u^2 really matters only when θ is sufficiently high that a low-k equilibrium exists.

Panel (d) of Figure 2 illustrates this point. We maintain θ at 0.40 as in Panel (b), but now set the ratio $\sigma_\varepsilon^2 / \sigma_u^2$ to 25 instead of 250. This leads to an increase in the (unique) equilibrium value of k from 0.04 to 0.35. Intuitively, with more noise in the rate-setting process, the market reacts less sensitively to a given change in rates, which allows the Fed to adjust rates more elastically in response to changes in its target.

C. The Time-Consistency Problem

A central property of our model is that there is a time-consistency problem: the Fed would do better in terms of optimizing its own objective function if it were able to commit to behaving as if it had a lower value of θ than it actually does. This is because while the Fed is always tempted to move the funds rate gradually so as to reduce bond-market volatility, this desire is frustrated in equilibrium: the more gradually the Fed acts, the more the market responds to any given change in rates. Thus, all that is left is a funds rate that is further from target on average. This point comes through most transparently if we consider the limiting case in which there is no noise in the Fed's rate-setting process. In this case, once we impose the rational-expectations assumption that $\kappa = k$, the Fed's attempts at smoothing have no effect at all on bond-market volatility in equilibrium:

$$\left(\Delta i_t^\infty\right)^2 = \left(\frac{k}{\kappa}\varepsilon_t\right)^2 = \varepsilon_t^2 \tag{14}$$

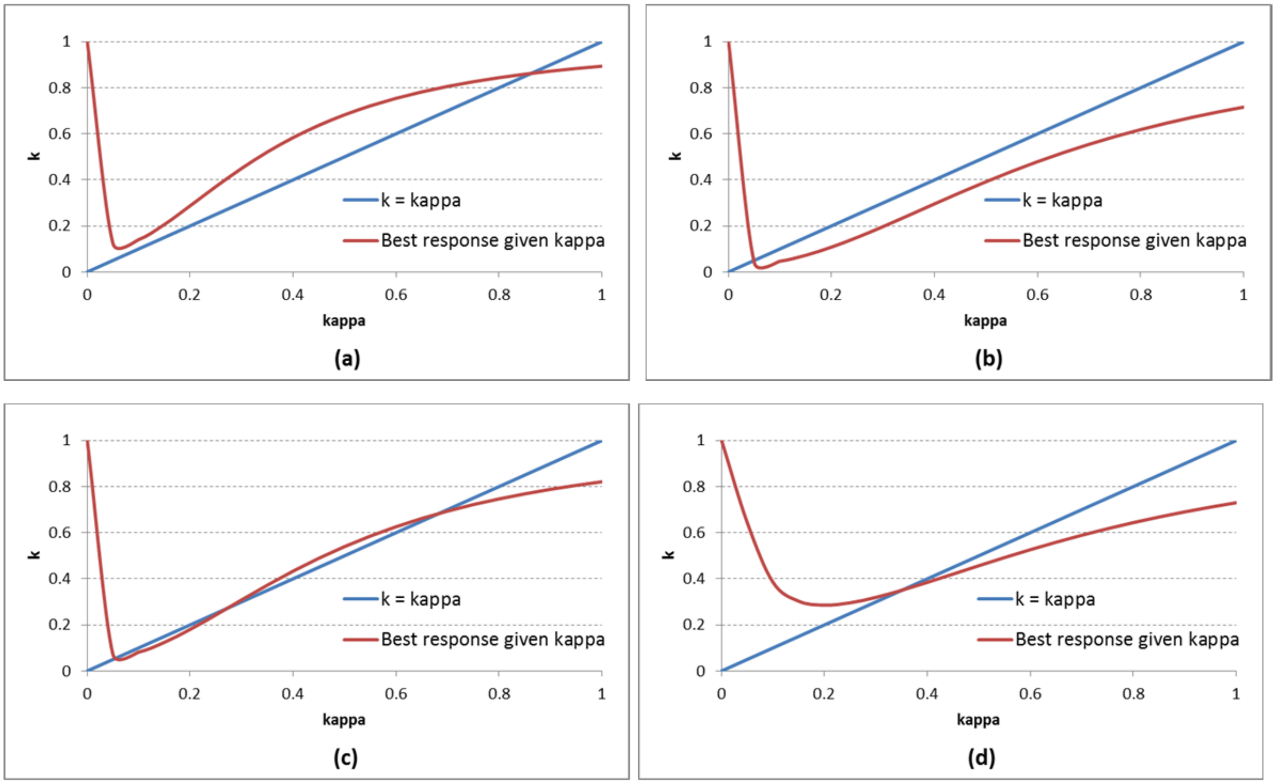

Figure 2. Behavior of the equilibrium for different parameter values. Panel (a): $\theta = 0.12$, $\sigma^2_{\varepsilon}/ \sigma^2_{u} = 250$.Panel (b): $\theta = 0.4$, $\sigma^2_{\varepsilon}/ \sigma^2_{u} = 250$. Panel (c): $\theta = 0.22$, $\sigma^2_{\varepsilon}/ \sigma^2_{u} = 250$. Panel (d): $\theta = 0.4$, $\sigma^2_{\varepsilon}/ \sigma^2_{u} = 25$.

Thus, the value of the Fed's loss function is

$$L_t = \left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^{\infty}\right)^2 = \left(\left(1-k\right)\varepsilon_t\right)^2 + \theta\varepsilon_t^2, \tag{15}$$

which is decreasing in k for $k < 1$.

To the extent that the target rate i_t^* is non-verifiable private information, it is hard to think of a contracting technology that can readily implement the first-best outcome under commitment: how does one write an enforceable rule that says that the Fed must always react fully to its private information? Thus, discretionary monetary policy will inevitably be associated with some degree of inefficiency. However, even in the absence of a binding commitment technology, there may still be scope for improving on the fully discretionary outcome. One possible approach follows in the spirit of Rogoff (1985), who argues that society should appoint a central banker who is more hawkish on inflation than is society itself. The analogy in the current context is that society should aim to appoint a central banker who cares less about financial-market volatility (i.e., has a lower value of θ) than society as a whole. Or said a little differently, society—and the central bank itself—should seek to foster an institutional culture and set of norms that discourages members of the monetary policy committee from being overly attentive to market-volatility considerations.

To see this, consider the problem of a social planner choosing a central banker whose concern about financial market volatility is given by θ_c. This central banker will implement the rational expectations adjustment rule $k(\theta_c)$, where k is given by equation (12), replacing θ with θ_c. The planner's ex ante problem is then to pick θ_c to minimize its ex ante loss, recognizing that its own concern about financial-market volatility is given by θ:

$$E_{\varepsilon_t}\left[\left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^\infty\right)^2\right] = E\left[\left(\left(1-k\left(\theta_c\right)\right)\varepsilon_t\right)^2 + \theta\varepsilon_t^2\right] = \left(\left(1-k\left(\theta_c\right)\right)^2 + \theta\right)\sigma_\varepsilon^2. \qquad (16)$$

Since $k(\theta_c)$ is decreasing in θ_c, the ex ante loss is minimized by picking $\theta_c = 0$, as the following proposition states.

Proposition 2: In the absence of rate-setting noise, it is ex ante optimal to appoint a central banker with $\theta_c = 0$ ***so that*** $k(\theta_c) = 1$.

In the alternative case where σ_u^2 is small but non-zero, a mild generalization of Proposition 2 obtains. While it is no longer true that the Fed would like to commit to behaving as if θ was exactly equal to zero, it would still like to commit to behaving as if θ was very close to zero and much smaller than its actual value. For example, for the parameter values in Panel D of Figure 2, where the Fed's actual θ = 0.4, it would minimize its loss function if it were able to commit to $\theta_c = 0.04$. In the Appendix, we provide a fuller treatment of the optimal solution under commitment in the case of non-zero noise.

D. Does Forward Guidance Help or Hurt?

Because the Fed's target rate i_t^* is non-verifiable private information, it is impossible to write a contract that effectively commits the Fed to rapid adjustment of the funds rate in the direction of i_t^*. So, as noted above, the first-best is generally unattainable. But this raises the question of whether there are other more easily enforced contracts that might be of some help in addressing the time-consistency problem. In this section, we ask whether something along the lines of forward guidance could fill this role. By "forward guidance", we mean an arrangement whereby the Fed announces at time t a value of the funds rate that it expects to prevail at some point in the future. Since both the announcement itself and the future realization of the funds rate are both publicly observable, market participants can easily tell ex post whether the Fed has honored its guidance, and one might imagine that it will suffer a reputational penalty if it fails to do so. In this sense, a "guidance contract" is more enforceable than a "don't-smooth-with-respect-to-private-information contract".

It is well understood that a semi-binding commitment to adhere to a pre-announced path for the short rate can be of value when the central bank is stuck at the zero lower bound (ZLB); see e.g., Krugman (1998), Eggertsson and Woodford (2003), and Woodford (2012). We do not take issue with this observation here. Rather, we ask a different question, namely whether guidance can also be useful away from the ZLB, as a device for addressing the central bank's tendency to adjust rates too gradually. It turns out the answer is no. Forward guidance can never help with the time-consistency problem we have identified, and attempts to use guidance away from the ZLB can potentially make matters strictly worse.

We develop the argument in two steps, working backwards. First, we assume that it is time t, and the Fed comes into the period having already announced a specific future rate path at time t-1. We then ask how the existence of this guidance influences the smoothing incentives analyzed above, and what impact it therefore has on the Fed's choice of i_t. Having done so, our second step is to fold back to time t-1 and ask what the optimal guidance announcement looks like and whether the existence of the guidance technology increases or decreases the Fed's expected utility from an ex ante perspective.

To be more precise, suppose we are in steady state at time t-1. We allow the Fed to publicly announce a value i_{t-1}^{f} as its guidance regarding i_t, the funds rate that will prevail one period later, at time t. To make the guidance partially credible, we assume that once it has been announced, the Fed bears a reputational cost of deviating from the guidance so that its augmented loss function at time t is now given by:

$$L_t = \left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^{\infty}\right)^2 + \gamma\left(i_{t-1}^{f} - i_t\right)^2 . \tag{17}$$

When it arrives at time t, the Fed takes the reputational penalty γ as exogenously fixed, but when we fold back to analyze guidance from an ex-ante perspective, it is more natural to think of γ as a choice variable for the Fed at time t-1. That is, the more emphatically it talks about its guidance ex ante or the more reputational chips it puts on the table, the greater will be the penalty if it deviates from the guidance ex post.

Again, our approach is to solve the Fed's problem by backwards induction, starting at time t and taking the forward guidance i_{t-1}^{f} as given. To keep things simple, we assume that there is no noise in the Fed's adjustment rule (σ_u^2 = 0). Recall from equation (13) above that in this case, there only exists a stable non-degenerate (i.e. positive k) equilibrium when θ < ¼, a condition that we assume to be satisfied in what follows. Moreover, in such an equilibrium, we have ½ < k < 1.

Given the more complicated nature of the Fed's objective function, its optimal choice as to how much to adjust the short rate no longer depends on just its private information ε_t; it also depends on the previously-established forward guidance i_{t-1}^{f}. To allow for a general treatment, we posit that this adjustment can be described by some potentially non-linear function of the two variables:

$$i_t = i_{t-1} + f\left(\varepsilon_t; i_{t-1}^{f}\right). \tag{18}$$

Moreover, the market conjectures that the Fed is following a potentially non-linear rule:

$$i_t = i_{t-1} + \phi\left(\varepsilon_t; i_{t-1}^{f}\right). \tag{19}$$

As before, given its conjecture about the Fed's adjustment rule, the market tries to back out i_t^* given i_t. Thus, the market's estimate of i_t^* satisfies

$$\begin{aligned} E\left[i_t^* \mid i_t\right] &= i_{t-1} + \phi^{-1}\left(i_t - i_{t-1}; i_{t-1}^{f}\right) \\ &= i_{t-1} + \phi^{-1}\left(f\left(\varepsilon_t, i_{t-1}^{f}\right); i_{t-1}^{f}\right). \end{aligned} \tag{20}$$

Based on the form of the Fed's adjustment rule and the fact that $i_t^{\infty} = E_t\left[i_t^*\right]$, the Fed's loss function can be written as:

$$\begin{aligned} &\left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^{\infty}\right)^2 + \gamma\left(i_{t-1}^{f} - i_t\right)^2 \\ &= \left(\varepsilon_t - f\left(\varepsilon_t; i_{t-1}^{f}\right)\right)^2 + \theta\left(\phi^{-1}\left(f\left(\varepsilon_t; i_{t-1}^{f}\right); i_{t-1}^{f}\right)\right)^2 + \gamma\left(i_{t-1}^{f} - i_{t-1} - f\left(\varepsilon_t; i_{t-1}^{f}\right)\right)^2. \end{aligned} \tag{21}$$

The Fed minimizes this loss function by choosing over the set of possible functions $f(\cdot;\cdot)$, taking the market's conjecture about its behavior $\phi(\cdot;\cdot)$ as fixed. In the Appendix, we use a calculus-of-variations type of argument to establish that, in a rational-expectations equilibrium in which $f(\cdot;\cdot) = \phi(\cdot;\cdot)$, the Fed's adjustment behavior is given by:

$$i_t = i_{t-1} + k\varepsilon_t + \frac{\gamma}{1+\gamma}\left(i_{t-1}^f - i_{t-1}\right), \tag{22}$$

where the partial-adjustment coefficient k now satisfies:

$$k = \left(1+\gamma\right)k^2 + \theta. \tag{23}$$

Two features of the Fed's modified adjustment rule are worth noting. First, as can be seen in (22), the Fed moves i_t more strongly in the direction of its previously-announced guidance i_{t-1}^f as γ increases, which is intuitive. Second, as (23) indicates, the presence of guidance also alters the way that the Fed incorporates its new time-*t* private information into the funds rate i_t. Given that any equilibrium necessarily involves $k > ½$, we can show that $\partial k / \partial \gamma < 0$. In other words, the more the Fed cares about deviating from its forward guidance, the less it responds to its new private information ε_t. In this sense, the presence of semi-binding forward guidance has an effect similar to that of increasing θ: both tend to promote gradualism with respect to the incorporation of new private information. The reason is straightforward: by definition, when guidance is set at time *t*-1, the Fed does not yet know the realization of ε_t. So ε_t cannot be impounded in the *t*-1 guidance, and anything that encourages the time-*t* rate to hew closely to the *t*-1 guidance must also discourage the incorporation of ε_t into the time-*t* rate. This is the key insight for why guidance is not helpful away from the ZLB.

To make this point more formally, let us now fold back to time *t*-1 and ask two questions. First, assuming that $\gamma > 0$, what is the optimal guidance announcement i_{t-1}^f for the Fed to make? And second, supposing that the Fed can choose the intensity of its guidance γ—i.e., it can choose how many reputational chips to put on the table when it makes the announcement—what is the optimal value of γ? In the Appendix, we demonstrate the following:

Proposition 3: Suppose we are in steady state at time t-1, with $i_{t-1} = i^*_{t-1}$. If the Fed takes $\gamma > 0$ as fixed, its optimal choice of forward guidance at time t-1 is to announce $i^f_{t-1} = i_{t-1}$. Moreover, given this announcement policy, the value of the Fed's objective function is decreasing in γ, so if it can choose, it is better off foregoing guidance altogether, i.e. setting $\gamma = 0$.

The first part of the proposition reflects the fact that because the Fed is already at its target at time *t*-1 and has no information about how that target will further evolve at time *t*, the best it can do is to set i^f_{t-1} at its current value of i_{t-1}. This guidance carries no incremental information to market participants at *t*-1 above and beyond what they can already deduce from observing the contemporaneous value of the funds rate. Nevertheless, even though it is uninformative, the guidance has an effect on rate-setting at time *t*. As equation (23) highlights, a desire to not renege on its guidance leads the Fed to underreact by more to its new time-*t* private information. Since this is strictly a bad thing, the Fed is better off not putting guidance in place to begin with.

This negative result about the value of forward guidance in Proposition 3 should be qualified in two ways. First, as we have already emphasized, this result only speaks to the desirability of guidance away from the ZLB; there is nothing here that contradicts Krugman (1998), Eggertsson and Woodford (2003), and Woodford (2012), who make the case for using a relatively strong form of guidance (i.e., with a significantly positive value of γ) when the economy is stuck at the ZLB.

Second, and more subtly, an advocate of using forward guidance away from the ZLB might argue that if it is done with a light enough touch, it can usefully help to communicate the Fed's private information to the market, without reducing the Fed's future room to maneuver. In particular, if the Fed is careful to make it clear that its guidance embodies no attempt at binding commitment whatsoever (i.e., that $\gamma = 0$), the guidance may be helpful purely as an informational device and, according to Proposition 3, can't hurt.

The idea that guidance can play a useful informational role strikes us as perfectly reasonable, even though it does not emerge in our model. Because the Fed's private information is uni-dimensional here, and because it is already fully revealed in equilibrium via the contemporaneous funds rate, guidance cannot serve to transmit any further information. Nevertheless, it is quite plausible that, in a richer model, things would be different, and guidance could be incrementally informative. On the other hand, our model also suggests that care should be taken when claiming that light-touch (γ = 0) guidance has no negative effects in terms of increasing the equilibrium degree of gradualism in the funds rate.

This point is perhaps most easily seen by considering the region of the parameter space where there are multiple equilibria. In this region, anything that influences market beliefs can have self-fulfilling effects. So for example, suppose that the Fed puts in place a policy of purely informational forward guidance, and FOMC members unanimously agree that the guidance does not in any way represent a commitment on their parts—that is, they all plan to behave as if $\gamma = 0$ and indeed manage to follow through on this plan ex post. Still, if we are in the multiple-equilibrium range, and the guidance leads some market participants to think in terms of a higher value of γ and therefore expect smaller deviations from the announced path of rates, this belief can be self-validating and can lead to an undesirable outcome where the Fed winds up moving more gradually in equilibrium, thereby lowering its utility. And again, this effect arises even if all FOMC members actually behave as if $\gamma = 0$ and attach no weight to keeping rates in line with previously-issued guidance. Thus, at least in the narrow context of our model, there would appear to be some potential downside associated with even the mildest forms of forward guidance once the economy is away from the ZLB.

E. Term Premium Shocks

We next enrich the model in another direction, so as to consider how the Fed behaves when financial-market conditions are not purely a function of the expected path of interest rates. Specifically, suppose that the infinite-horizon forward rate consists of both the expected future short rate and an exogenous term premium component r_t:

$$i_t^{\infty} = E_t\left[i_t^*\right] + r_t. \tag{24}$$

The term premium is assumed to be common information, observed simultaneously by market participants and the Fed. We allow the term premium to follow an arbitrary process and let η_t denote the innovation in the term premium:

$$\eta_t = r_t - E_{t-1}\left[r_t\right]. \tag{25}$$

The solution method in this case is similar to that used for forward guidance in Section III.D. Specifically, we again assume that there is no noise in the Fed's rate-setting rule (σ_u^2 = 0) but that the rule can be an arbitrary non-linear function of both the new private information ε_t that the Fed learns at time *t* as well as the term premium shock η_t, which is publicly observable:

$$i_t = i_{t-1} + f\left(\varepsilon_t;\eta_t\right). \tag{26}$$

The market again conjectures that the Fed is following a non-linear rule:

$$i_t = i_{t-1} + \phi\left(\varepsilon_t;\eta_t\right). \tag{27}$$

As before, given its conjecture about the Fed's adjustment rule, the market tries to back out i_t^*. Thus, the market's assessment of i_t^* given i_t satisfies:

$$E\left[i_t^* \mid i_t\right] = i_{t-1} + \phi^{-1}\left(f\left(\varepsilon_t;\eta_t\right);\eta_t\right). \tag{28}$$

Thus, the Fed picks the adjustment function $f\left(\cdot;\cdot\right)$ to minimize:

$$\left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^\infty\right)^2 = \left(\varepsilon_t - f\left(\varepsilon_t;\eta_t\right)\right)^2 + \theta\left(\phi^{-1}\left(f\left(\varepsilon_t;\eta_t\right);\eta_t\right)+\eta_t\right)^2. \tag{29}$$

In the Appendix, we show that the solution to (29), combined with the rational-expectations condition that $f\left(\cdot;\cdot\right) = \phi(\cdot;\cdot)$, leads to a Fed adjustment rule given by:

$$i_t = i_{t-1} + k_\varepsilon \varepsilon_t + k_\eta \eta_t, \tag{30}$$

where

$$k_\varepsilon = k_\varepsilon^2 + \theta \qquad \text{and}$$
$$k_\eta = -\frac{\theta}{k_\varepsilon}. \tag{31}$$

The following proposition summarizes the key properties of the equilibrium.

Proposition 4: The Fed acts to offset term premium shocks, lowering the funds rate when the term premium shock is positive and raising it when the term premium shock is negative. It does so even though term premium shocks are publically observable and in spite of the fact that its efforts to reduce the volatility associated with term premium shocks are completely fruitless in equilibrium. As the Fed's concern with bond-market volatility θ increases, k_ε falls and k_η increases in absolute magnitude. Thus, when it cares more about the bond market, the Fed reacts more gradually to changes in its private information about its target rate but more aggressively to changes in term premiums.

A first observation about the equilibrium is that the Fed's response to its own private information, k_{ε}, is exactly the same as it was in the absence of term premium shocks (assuming no rate-setting noise) as derived in Section III.B.1. One might then be tempted to think that nothing changes when we introduce term premium shocks, because they are publicly observable. Proposition 4 shows that, strikingly, this is not the case. Short rates do in fact respond to movements in the term premium.

What explains this result? Essentially, when the term premium spikes up, the Fed is unhappy about the prospective increase in the volatility of long rates. So even if its private information about i_t^* has in fact not changed, it would like to make the market think it has become more dovish so as to offset the rise in the term premium. Therefore, it cuts the short rate in an effort to create this impression. Again, in equilibrium, this attempt to fool the market is not successful, but taking the market's conjectures at any point in time as fixed, the Fed is always tempted to try.

One way to see why equilibrium must involve the Fed reacting to term premium shocks is to think about what happens if we try to sustain an equilibrium where it doesn't—that is, if we try to sustain an equilibrium in which $k_{\eta} = 0$. In such a hypothetical equilibrium, equation (30) tells us that when the market sees any movement in the funds rate, it attributes that movement entirely to changes in the Fed's private information ε_t about its target rate. But if this is the case, then the Fed can indeed offset movements in term premiums by changing the short rate, thereby contradicting the assumption that $k_{\eta} = 0$. Hence, $k_{\eta} = 0$ cannot be an equilibrium.

A second key feature of the equilibrium is that the absolute magnitude of k_{η} becomes larger as θ rises and as k_{ε} becomes smaller: when it cares more about bond-market volatility, the Fed's responsiveness to term premium shocks becomes more aggressive even as its adjustment to new private information becomes more gradual. In particular, because we are restricting ourselves to the region of the parameter space where the simple no-noise model yields a non-degenerate equilibrium for k_{ε}, this means (from equation (13) above) that we must have $0 < \theta < ¼$. As θ moves from the lower to the upper end of this range, k_{ε} declines monotonically from 1 to ½, and k_{η} increases in absolute magnitude from 0 to –½.

This latter part of the proposition is especially useful, as it yields a sharp testable empirical implication. As noted earlier, Campbell, Pflueger, and Viceira (2015) have shown that the Fed's behavior has become significantly more inertial in recent years. If we think of this change as reflecting a lower value of k_ε, we might be tempted to use the logic of the model to claim that the lower k_ε is the result of the Fed placing increasing weight over time on the bond market, i.e. having a higher value of θ than it used to. While the evidence on financial-market mentions in the FOMC transcripts that we plotted in Figure 1 is loosely consistent with this hypothesis, it is obviously far from being a decisive test. However with Proposition 4 in hand, if we want to attribute a decline in k_ε to an increase in θ, then we also have to confront the additional prediction that we ought to observe the Fed responding more forcefully over time to term premium shocks. That is, the absolute value of k_η must have gone up. If this is not the case, it would represent a rejection of the hypothesis.

Our earlier results about the value of commitment generalize straightforwardly to the case where there are term premium shocks. Consider, as before, the problem of a social planner appointing a central banker whose concern about financial market volatility is given by θ_c. This central banker will implement the adjustment rules characterized by $k_\varepsilon(\theta_c)$ and $k_\eta(\theta_c)$, which are given by equation (31), replacing θ with θ_c. Thus, the planner's ex ante problem is to pick θ_c to minimize:

$$E\left[\left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^\infty\right)^2\right] = E\left[\left(\left(1 - k_\varepsilon(\theta_c)\right)\varepsilon_t - k_\eta(\theta_c)\eta_t\right)^2 + \theta\left(\varepsilon_t + \eta_t\right)^2\right]. \tag{32}$$

As can be seen from (32), the planner's optimum is attained when $k_\varepsilon(\theta_c) = 1$ and $k_\eta(\theta_c) = 0$, which, according to (31), can be implemented by setting $\theta_c = 0$. Thus we have:

Proposition 5: In the extended version of the model with term premium shocks, it is again optimal to appoint a central banker with $\theta_c = 0$ **.**

As before, in the rational-expectations equilibrium with term premium shocks, the Fed cannot influence the volatility of the infinite horizon-forward rate. Thus, if the Fed moves gradually, it just ends up with a funds rate that is further from target on average, with no benefit in terms of reduced volatility. Therefore, it is optimal to commit to not moving gradually by setting $\theta_c = 0$.

IV. Dynamic Model

In this section, we introduce a dynamic extension of the model. The dynamic model allows us to study the behavior of finite-horizon yields and forward rates and to show that our basic point about the existence of a time-consistency problem is robust in a dynamic setting.

A. Dynamic Setup

To study dynamics, we need to consider the Fed's behavior when the economy is not in steady state at time *t*-1. Suppose we enter period *t* with a pre-existing gap between the time *t*-1 target rate and the time *t*-1 federal funds rate of $X_{t-1} \equiv i^*_{t-1} - i_{t-1}$. If we continue to assume that there is no noise in the Fed's rate-setting process ($\sigma_u^2 = 0$), we know that in equilibrium, the Fed's actions at time *t*-1 fully reveal X_{t-1}. Thus, the market is fully aware of X_{t-1} by the beginning of time *t*.

We assume for simplicity that the log expectations hypothesis holds, so that we can ignore the term premium. We further assume that the Fed's time *t* adjustment rule is of the form:

$$i_t = i_{t-1} + k_t(i_t^* - i_{t-1}) = i_{t-1} + k_t\left(X_{t-1} + \varepsilon_t\right). \tag{33}$$

That is, the Fed adjusts gradually to the gap $(i_t^* - i_{t-1})$, which implies that it treats its newly arrived private information ε_t and its pre-existing public information X_{t-1} symmetrically. This restriction, which is important for the analysis that follows, is a substantive one. In our model, since X_{t-1} is already known by the market at time *t*, it is already incorporated into bond yields. Thus, if we gave it the freedom to adjust differentially to X_{t-1} and ε_t, the Fed would choose to fully adjust the funds rate to X_{t-1}, even as it partially adjusted to ε_t, because it should be able to do the former without inducing additional bond-market volatility.

In practice, there may exist a variety of reasons for the Fed to treat ε_t and X_{t-1} symmetrically. For example, it may be a natural heuristic, because it is presumably easier for individual FOMC members to think in terms of the overall gap $(i_t^* - i_{t-1})$ between their target and where rates currently stand, as opposed to carrying around a decomposition of this gap. Moreover, if we were to re-introduce rate-setting noise into the model (i.e., if we allowed for $\sigma_u^2 > 0$), it would no longer be true that X_{t-1} is fully revealed to the market by time *t*. In this case, the Fed would still want to under-adjust somewhat to X_{t-1}. And even if the optimal degree of under-adjustment to X_{t-1} and ε_t were not the same, this desire to under-adjust to both might reinforce a rule-of-thumb heuristic where FOMC members tend to treat the two in the same way.

To further simplify the problem, we also assume the following timing convention. First, based on its knowledge of X_{t-1}, the Fed decides on the value of the adjustment parameter k_t it will use for the time *t* FOMC meeting. After making this decision, it deliberates further, and in so doing discovers the committee's consensus value of ε_t. Thus, the Fed picks k_t, taking X_{t-1} as given, but before knowing the realization of ε_t. This timing convention is purely a technical trick that makes the problem more tractable without really changing anything of economic substance. Without it, the Fed's adjustment rule would turn out to depend on the realization of the product $X_{t-1}\varepsilon_t$. With the timing trick, what matters instead is the *expectation* of the product $X_{t-1}\varepsilon_t$, which is zero. This simplification maintains the linearity of the Fed's optimal adjustment rule. We should emphasize that even with this somewhat strained intra-meeting timing, the Fed still behaves on a discretionary basis from one meeting to the next. Thus, while it agrees to a value of k_t in the first part of the time-*t* meeting, it has no ability to bind itself to a value of k_t across meetings. Hence the basic commitment problem remains.

Given these assumptions, the rational expectations equilibrium is an intuitive generalization of the one derived in Section III.B above. The speed of adjustment k_t is now time-varying and satisfies:

$$k_t = k_t^2 + \theta \frac{\sigma_\varepsilon^2}{X_{t-1}^2 + \sigma_\varepsilon^2}. \tag{34}$$

When we are in steady state at time *t*-1, with $X_{t-1} = 0$, this expression reduces to $k_t = k_t^2 + \theta$, which is the same result we had in Section III.B.1. When we are not in steady state, an increase in X_{t-1}^2 has an effect that is isomorphic to reducing the value of θ; that is, it increases the equilibrium value of k_t. Thus the model has the natural property that the speed of adjustment is faster the larger is the pre-existing gap between the funds rate and the target. In other words, the Fed behaves in a less gradual fashion when it comes into period *t* with a lot of ground to make up.

B. Behavior of Finite-Horizon Forward Rates

Now that we know how the Fed behaves on a period-by-period basis, we can study the dynamic path of the federal funds rate, which is characterized in the following proposition.

Proposition 6: Suppose we start in steady state with $X_{t-1} = 0$***. Then, the subsequent evolution of the federal funds rate is given by:*** $i_{t+n} = \sum_{j=0}^{n} \left(1 - \left(\prod_{h=j}^{n} (1 - k_{t+h}) \right) \right) \varepsilon_{t+j}$.

The proposition states that i_{t+n} is effectively the accumulation of a series of shocks to the Fed's target rate ε_{t+j}. Each of these shocks is gradually incorporated into the funds rate over time, with a trajectory of adjustment that depends on the realization of the path of k_{t+h}.

Because we assume there is no term premium, the *n*-period forward rate at time *t* is given by $i_t^n = E_t\left[i_{t+n}\right]$. We cannot solve for these finite-horizon forward rates in closed form, because the future adjustment speeds k_{t+h} are stochastic and depend on the realization of the sequence of ε_{t+j}. Nonetheless, the implicit characterization in Proposition 6 is sufficient to establish the basic properties of the equilibrium, which are given by the following corollary.

Corollary: Increased smoothing (lower values of k_{t+h}***) reduces the volatility of the finite horizon forward rate*** i_t^n***. The effect on volatility is smaller for longer horizons, i.e., for larger values of n.***

In contrast to the case of the infinite-horizon forward rate studied above, here the Fed's gradual adjustment to its target rate does actually lower the equilibrium volatility of finite-horizon forward rates. So in this setting, it is no longer accurate to say that the Fed's attempts to mitigate bond-market volatility are completely frustrated, if the volatility that we are referring to is that of a finite-horizon rate. To see why, imagine that the target rate i_t^* changes instantaneously from 1% to 5%, but the Fed moves so slowly in response this change that the funds rate only increases by 50 basis points per year. Then, even if the large increase in i_t^* is fully revealed in equilibrium, the two-year forward rate will only go up by 100 basis points—so gradualism does indeed matter for this rate. Of course, the longer the horizon, the more of the Fed's ultimate adjustment the forward rate will reflect. In the limiting case of the infinite-horizon forward rate, the market knows the funds rate will eventually reach 5%, so the infinite-horizon forward rate jumps to 5% immediately, irrespective of the degree of gradualism.

C. Is There Still a Time-Consistency Problem?

When we focused on the infinite-horizon forward rate in Section III.C, the time-consistency problem was very stark: the Fed's efforts to mitigate bond-market volatility were shown to be completely unsuccessful in equilibrium, and all the Fed got for these efforts was a funds rate that at any point in time was further away from its dual-mandate target. However, given that we now see that Fed gradualism does have some equilibrium effect on finite-horizon forward rates, it is natural to ask whether there is still a time-consistency problem if the Fed's objective function depends on the volatility of finite-horizon rates as opposed to the volatility of infinite-horizon rates. As we demonstrate below, the answer is yes.

As we saw in the previous section, the behavior of finite-horizon forward rates depends on the market's conjectures about the path of future speeds of adjustment k_{t+h}. In turn, the optimal speeds of adjustment chosen by the Fed depend on the finite-horizon forward rate that appears in its objective function. Solving for the equilibrium recognizing this feedback loop is quite complex. For tractability, we simplify the problem by assuming that the Fed's objective function depends on an approximation of the finite-horizon forward rate that we call the α-pseudorate. The α-pseudorate is a weighted average of the current federal funds rate i_t and the infinite-horizon forward rate i_t^∞:

$$i_t^{psu} = \alpha i_t + (1-\alpha) i_t^\infty . \tag{35}$$

In the equilibrium we solve for, it will be the case that the n-year forward rate can be closely approximated by the α-pseudorate for the appropriately-chosen value of α. This follows from the fact that there is only one factor, ε_t, in the model of interest rates we have written down. Different finite-horizon forward rates load differently on ε_t, but given our assumptions all volatility in any finite-horizon forward rate is ultimately driven by ε_t.[10]

The following proposition shows that the basic time-consistency problem carries over to this setting, with discretionary behavior leading to excessive gradualism.

Proposition 7: Suppose that $\theta > 0$ and that the Fed cares about the volatility of the α-pseudorate, i_t^{psu}. Then, the Fed would like to commit itself to $\theta_c < \theta$. The Fed's adjustment speed at time t, k_t, is greater if the Fed can commit itself than if the Fed behaves on a discretionary basis.

Under discretion, the Fed has two motives when it moves gradually. First, it aims to reduce the volatility of the component of the α-pseudorate related to the current short rate i_t. Second, it hopes to reduce the volatility of the component of the α-pseudorate related to the infinite-horizon forward rate i_t^{∞} by fooling the market about its private information ε_t. The first goal can be attained in rational-expectations equilibrium, but the second cannot—as in the static model, the Fed cannot fool the market about its private information in equilibrium.

Under commitment, only one of these motives for moving gradually remains, and thus it is less appealing to move gradually than it would be under discretion. Thus, if society is appointing a central banker whose concern about financial market volatility is given by θ_c, it would like to appoint one with $\theta_c < \theta$. However, unlike the case where the social planner cares about the volatility of the infinite-horizon forward rate, here the planner no longer wants to set $\theta_c = 0$. Because moving gradually does partially succeed in reducing the volatility of the α-pseudorate, the Fed's efforts are not completely frustrated. Thus, it is indeed optimal for the Fed to move somewhat gradually, so that the optimal θ_c is greater than zero. However, it is not optimal for the Fed to move as gradually as it would under discretion.

---

[10] However, assuming that the Fed's objective depends on the n-year forward rate rather than the α-pseudorate changes the nature of the equilibrium itself. In other words, the speeds of adjustment k_{t+h} in the rational-expectations equilibrium will be different depending on which assumption is used.

V. Conclusion

Rather than restating our findings, we conclude with a brief caveat. We have argued that in our setting, it can be valuable for a central bank to develop an institutional culture and set of norms such that a concern with bond-market volatility does not play an outsized role in policy deliberations. In other words, it can be useful for monetary policymakers to build a reputation for not caring too much about the bond market. This argument rests on a comparison of equilibrium outcomes with different values of the parameter θ_c, which measures the appointed central banker's concern about financial-market volatility. But crucially, an implicit assumption in making this comparison is that in any given equilibrium, the market has come to fully know the true value of θ_c. We have not addressed the more difficult question of out-of-equilibrium dynamics: how is it that the market learns about θ_c, either from the central bank's observed behavior, or through other forms of communication?

For this reason, our model offers no guidance on the best way to make the transition to a less bond-market-sensitive equilibrium, one in which there is less inertia in the policy rate, and the market eventually comes to understand that θ_c has declined. At any point in time, taking market conjectures as fixed, a "cold turkey" approach that involves an unexpectedly sharp adjustment in the path of the policy rate (relative to prevailing expectations of gradualism) is likely to cause significant market volatility, and perhaps some collateral costs for the real economy. To be clear, nothing in our analysis should be taken as advocating such an approach. In particular, we do not believe our model has much advice to give with respect to the near-term question of how rapidly the Fed should lift rates in the months following its liftoff from the zero lower bound. At this horizon, market conjectures regarding the equilibrium degree of gradualism are likely to be relatively fixed, and the costs of changing behavior suddenly could potentially be quite high. To the extent that an institution-building effort of the sort we have in mind is at all useful, it may be something that is better undertaken over a longer horizon, or at a time when the economy is in a less fragile state.

References

Barro, Robert J., and David B. Gordon, 1983. "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics* 12, 101-121.

Bernanke, Ben, 2004. "Gradualism," Remarks at an economics luncheon co-sponsored by the Federal Reserve Bank of San Francisco and the University of Washington, Seattle, Washington.

Brainard, William, 1967. "Uncertainty and the Effectiveness of Policy," *American Economic Review* 57, 411-425.

Coibion, Olivier, and Yuriy Gorodnichenko, 2012. "Why Are Target Interest Rate Changes So Persistent?" *American Economic Journal: Macroeconomics* 4, 126-62.

Campbell, John Y., Carolin Pflueger, and Luis Viceira, 2015. "Monetary Policy Drivers of Bond and Equity Risks," Harvard University working paper.

Eggertsson, Gauti B. and Michael Woodford, 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 139-211.

Farhi, Emmanuel, and Jean Tirole, 2012. "Collective Moral Hazard, Maturity Mismatch, and Systemic Bailouts," *American Economic Review* 102, 60-93.

Holmstrom, Bengt, 1999. "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies* 66, 169-182.

Krugman, Paul, 1998. It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 137-2006.

Kydland, Finn, and Edward Prescott, 1977. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85, 473-491.

Rogoff, Kenneth, 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target," *Quarterly Journal of Economics* 100, 1169-1189.

Rudebusch, Glenn D., 2002. "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia," *Journal of Monetary Economics* 49, 1161-1187.

Rudebusch, Glenn D., 2006. "Monetary Policy Inertia: Fact or Fiction?" *International Journal of Central Banking* 2, 85-135.

Sack, Brian, 1998. "Uncertainty, Learning, and Gradual Monetary Policy," FEDS Working Paper Series, number 98-34.

Stein, Jeremy C., 1989. "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *Quarterly Journal of Economics* 104, 655−669.

Stein, Jeremy C., 2014. "Challenges for Monetary Policy Communication," Remarks at the Money Marketeers of New York University, New York, May 6.

Taylor, John, 1993. "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214.

Woodford, Michael, 2003. "Optimal Interest-Rate Smoothing," *Review of Economic Studies* 70, 861-886.

Woodford, Michael, 2012. "Methods of Policy Accommodation at the Interest-Rate Lower Bound," *The Changing Policy Landscape: 2012 Jackson Hole Symposium*, Federal Reserve Bank of Kansas City.

A Proofs of Propositions

A.1 Proof of Proposition 1

The change in the infinite-horizon forward rate i_t^{∞} is given by $E[\varepsilon_t | i_t]$. The market observes

$$i_t = i_{t-1} + \kappa\varepsilon_t + u_t$$

so it has a noisy signal of ε_t given by

$$\widetilde{\varepsilon}_t = \frac{i_t - i_{t-1}}{\kappa} = \varepsilon_t + \frac{u_t}{\kappa},$$

Since both ε_t and u_t are normally distributed, the market's posterior over ε_t is given by

$$\frac{\kappa^2\tau_u}{\tau_\varepsilon + \kappa^2\tau_u}\widetilde{\varepsilon}_t = \frac{\kappa\tau_u}{\tau_\varepsilon + \kappa^2\tau_u}\left(i_t - i_{t-1}\right) = \chi\left(i_t - i_{t-1}\right).$$

Thus the Fed's loss function is given by

$$(1-k)^2\varepsilon_t^2 + \sigma_u^2 + \theta\chi^2\left(k^2\varepsilon_t^2 + \sigma_u^2\right).$$

Differentiating with respect to k yields the first order condition

$$k = \frac{1}{1+\theta\chi^2}.$$

In rational expectations, we have $\kappa = k$ so we have

$$k = \frac{\left(\tau_\varepsilon + k^2\tau_u\right)^2}{\left(\tau_\varepsilon + k^2\tau_u\right)^2 + \theta\left(k\tau_u\right)^2}. \tag{A.1}$$

Note that as $\tau_u \to \infty$ so that there is no noise, this collapses to

$$k = \frac{k^2}{k^2+\theta}. \tag{A.2}$$

A.2 Proof of Proposition 2

Society's ex ante problem is to choose a central banker with concern about market volatility θ_c. In the absence of noise, this central banker will implement the rational-expectations equilibrium $k(\theta_c)$ given by Equation (A.2) replacing θ with θ_c. In the absence of noise, society's ex ante loss function is given by

$$\left((1-k(\theta_c))^2 + \theta\right)\sigma_\varepsilon^2$$

which is minimized by setting $\theta_c = 0$ so that $k(\theta_c) = 1$.

In the presence of noise, the appointed central banker will implement the rational-expectations equilibrium $k(\theta_c)$ given by Equation (A.1) replacing θ with θ_c. Differentiating $k(\theta_c)$ with respect to θ_c yields

$$\frac{\partial k}{\partial\theta_c} = -\left(1 - \frac{2\theta_c k\tau_u^2\left(\tau_\varepsilon + k^2\tau_u\right)\left(k^2\tau_u - \tau_\varepsilon\right)}{\left(\left(\tau_\varepsilon + k^2\tau_u\right)^2 + \theta_c\left(k\tau_u\right)^2\right)^2}\right)^{-1}\frac{\left(\tau_\varepsilon + k^2\tau_u\right)^2\left(k\tau_u\right)^2}{\left(\left(\tau_\varepsilon + k^2\tau_u\right)^2 + \theta_c\left(k\tau_u\right)^2\right)^2} < 0.$$

In the presence of noise, society's ex ante loss function is given by

$$L = E\left[\left((1-k)\varepsilon_t + u_t\right)^2 + \theta\left(\chi\left(k\varepsilon + u_t\right)\right)^2\right] = (1-k)^2\sigma_\varepsilon^2 + \sigma_u^2 + \theta\chi^2\left(k^2\sigma_\varepsilon^2 + \sigma_u^2\right).$$

Substituting in for the rational-expectations definition of χ and differentiating with respect to θ_c gives the first order condition

$$\frac{\partial L}{\partial \theta_c} = \left[-2(1-k)\sigma_\varepsilon^2 + \frac{2k\theta\tau_u}{\left(\tau_\varepsilon + k^2\tau_u\right)^3}\left(k^2\tau_u + \tau_\varepsilon\right)\right]\frac{\partial k}{\partial \theta_c}$$

At the optimum θ_c in the presence of noise, the term in square brackets will be zero. When we discuss the optimal θ_c in the presence of noise in the main text, we numerically solve for the value of θ_c and therefore k that sets the term in square brackets to zero.

A.3 Proof of Proposition 3

We begin by solving the time t problem, taking the forward guidance i_{t-1}^f as given. Given its conjecture about the rule the Fed is following, $\phi\left(\varepsilon_t; i_{t-1}^f\right)$, the market's conjecture about ε_t is

$$\widetilde{\varepsilon_t} = \phi^{-1}\left(f\left(\varepsilon_t; i_{t-1}^f\right); i_{t-1}^f\right)$$

and the Fed's loss function is

$$\left(\varepsilon_t - f\left(\varepsilon_t; i_{t-1}^f\right)\right)^2 + \theta\left(\phi^{-1}\left(f\left(\varepsilon_t; i_{t-1}^f\right); i_{t-1}^f\right)\right)^2 + \gamma\left(i_{t-1}^f - i_{t-1} - f\left(\varepsilon_t; i_{t-1}^f\right)\right)^2.$$

Consider the effect on the value of the loss function of a small perturbation in the value of $f\left(\varepsilon_t; i_{t-1}^f\right)$, df. The effect of this perturbation is zero at the optimal $f(\cdot;\cdot)$ so we have

$$-\left(\varepsilon_t - f\left(\varepsilon_t; i_{t-1}^f\right)\right) + \theta\left(\phi^{-1}\left(f\left(\varepsilon_t; i_{t-1}^f\right); i_{t-1}^f\right)\right)\frac{\partial \phi^{-1}}{\partial i}\Big|_{i=f\left(\varepsilon_t; i_{t-1}^f\right)} - \gamma\left(i_{t-1}^f - i_{t-1} - f\left(\varepsilon_t; i_{t-1}^f\right)\right) = 0.$$

Since $\phi\left(\phi^{-1}(x)\right) = x$, we have

$$\frac{\partial \phi^{-1}}{\partial i}\Big|_{i=f\left(\varepsilon_t; i_{t-1}^f\right)} = \frac{1}{\frac{\partial \phi}{\partial i}\Big|_{\varepsilon=\phi^{-1}\left(f\left(\varepsilon_t; i_{t-1}^f\right); i_{t-1}^f\right)}}.$$

So the optimal $f(\cdot;\cdot)$ satisfies:

$$-\left(\varepsilon_t - f\left(\varepsilon_t; i_{t-1}^f\right)\right) + \theta\left(\phi^{-1}\left(f\left(\varepsilon_t; i_{t-1}^f\right); i_{t-1}^f\right)\right)\frac{1}{\frac{\partial \phi}{\partial i}\Big|_{\varepsilon=\phi^{-1}\left(f\left(\varepsilon_t; i_{t-1}^f\right); i_{t-1}^f\right)}} - \gamma\left(i_{t-1}^f - i_{t-1} - f\left(\varepsilon_t; i_{t-1}^f\right)\right) = 0.$$

Imposing rational expectations, we have $\phi = f$ so that this reduces to the differential equation

$$\frac{\partial f}{\partial \varepsilon}\Big|_{\varepsilon=\varepsilon_t}\left(\varepsilon_t - f\left(\varepsilon_t; i_{t-1}^f\right) + \gamma\left(i_{t-1}^f - i_{t-1} - f\left(\varepsilon_t; i_{t-1}^f\right)\right)\right) = \theta\varepsilon_t,$$

which the optimal $f(\cdot;\cdot)$ must satisfy.

Now conjecture that $f = k\varepsilon_t + c$. In this case the differential equation reduces to

$$k\left(\varepsilon_t - k\varepsilon_t - c + \gamma\left(i_{t-1}^f - i_{t-1} - k\varepsilon_t - c\right)\right) = \theta\varepsilon_t$$

or

$$k(1-k-\gamma k)\varepsilon_t - kc + k\gamma\left(i_{t-1}^f - i_{t-1} - c\right) = \theta\varepsilon_t$$

Matching coefficients implies that

$$\begin{aligned} k(1-k-\gamma k) &= \theta \quad \text{and} \\ -kc + k\gamma\left(i_{t-1}^f - i_{t-1} - c\right) &= 0. \end{aligned} \tag{A.3}$$

Given the rule $f\left(\varepsilon_t; i_{t-1}^f\right)$ the Fed will follow at time t, we can now solve for the optimal forward guidance at $t-1$. The loss function at time t is given by

$$\begin{aligned} & \left((1-k)\varepsilon_t - c\right)^2 + \theta\varepsilon_t^2 + \gamma\left(i_{t-1}^f - i_{t-1} - k\varepsilon_t - c\right)^2 \\ = & \left((1-k)\varepsilon_t - c\right)^2 + \theta\varepsilon_t^2 + \gamma\left(k\varepsilon_t + \frac{1}{\gamma}c\right)^2 \end{aligned}$$

which in expectation at $t-1$ is equal to

$$\left((1-k)^2 + \gamma k^2 + \theta\right)\sigma_\varepsilon^2 + \frac{\gamma}{1+\gamma}\left(i_{t-1}^f - i_{t-1}\right)^2. \tag{A.4}$$

This is minimized by setting $i_{t-1}^f = i_{t-1}$.

Finally we solve for the optimal γ. Differentiating (A.3) with respect to γ, we have

$$\frac{\partial k}{\partial \gamma} = \frac{-k^2}{2k(1+\gamma)-1} < 0$$

so long as $k > 1/2$, which it will be in the parameter space we study. Setting $i_{t-1}^f = i_{t-1}$ and differentiating the ex ante loss function (A.4) with respect to γ yields

$$\left(-2(1-k)\frac{\partial k}{\partial \gamma} + k^2\left(1 - \frac{2k\gamma}{2k(1+\gamma)-1}\right)\right)\sigma_\varepsilon^2 > 0$$

since $\partial k/\partial\gamma < 0$ and $k > 1/2$. Since the ex ante loss is increasing in γ, it is optimal to set $\gamma = 0$.

A.4 Proof of Proposition 4

Given its conjecture about the rule the Fed is following, $\phi(\varepsilon_t;\eta_t)$, the market's conjecture about ε_t is

$$\widetilde{\varepsilon_t} = \phi^{-1}\left(f(\varepsilon_t;\eta_t);\eta_t\right)$$

and the Fed's loss function is

$$\left(\varepsilon_t - f(\varepsilon_t;\eta_t)\right)^2 + \theta\left(\phi^{-1}\left(f(\varepsilon_t;\eta_t);\eta_t\right) + \eta_t\right)^2.$$

Consider the effect on the value of the loss function of a small perturbation in the value of $f(\varepsilon_t;\eta_t)$, df. The effect of this perturbation is zero at the optimal $f(\cdot;\cdot)$ so we have

$$-\left(\varepsilon_t - f(\varepsilon_t;\eta_t)\right) + \theta\left(\phi^{-1}\left(f(\varepsilon_t;\eta_t);\eta_t\right) + \eta_t\right)\frac{\partial\phi^{-1}}{\partial i}|_{i=f(\varepsilon_t;\eta_t)} = 0. \tag{A.5}$$

Since $\phi\left(\phi^{-1}(x)\right)=x$, we have

$$\frac{\partial\phi^{-1}}{\partial i}|_{i=f(\varepsilon_t;\eta_t)}=\frac{1}{\frac{\partial\phi}{\partial i}|_{\varepsilon=\phi^{-1}(f(\varepsilon_t;\eta_t);\eta_t)}}.$$

Substituting into (A.5) gives

$$-\left(\varepsilon_t-f\left(\varepsilon_t;\eta_t\right)\right)+\theta\left(\phi^{-1}\left(f\left(\varepsilon_t;\eta_t\right);\eta_t\right)+\eta_t\right)\frac{1}{\frac{\partial\phi}{\partial i}|_{\varepsilon=\phi^{-1}(f(\varepsilon_t;\eta_t);\eta_t)}}=0.$$

Imposing rational expectations, we have $\phi=f$ so that this reduces to the differential equation

$$\frac{\partial f}{\partial\varepsilon}|_{\varepsilon=\varepsilon_t}\left(\varepsilon_t-f\left(\varepsilon_t;\eta_t\right)\right)=\theta\left(\varepsilon_t+\eta_t\right),$$

which the optimal $f(\cdot;\cdot)$ must satisfy.

Now conjecture that $f=k_\varepsilon\varepsilon_t+c$. In this case the differential equation reduces to

$$k_\varepsilon\left(\varepsilon_t-k_\varepsilon\varepsilon_t-c\right)=\theta\left(\varepsilon_t+\eta_t\right)$$

or

$$k_\varepsilon\left(1-k_\varepsilon\right)\varepsilon_t-k_\varepsilon c=\theta\varepsilon_t+\theta\eta_t$$

Matching coefficients yields

$$\begin{aligned} k_\varepsilon\left(1-k_\varepsilon\right) &= \theta \quad \text{and} \\ c &= -\frac{\theta}{k_\varepsilon}\eta_t. \end{aligned} \tag{A.6}$$

Thus, we can write the optimal $f(\cdot;\cdot)$ as

$$f=k_\varepsilon\varepsilon_t+k_\eta\eta_t$$

where $k_\varepsilon\left(1-k_\varepsilon\right)=\theta$ and $k_\eta=-\theta/k_\varepsilon$.

A.5 Proof of Proposition 6

We begin by solving for the Fed's optimal rule when its loss function is over the infinite-horizon forward rate. The market observes

$$i_t=i_{t-1}+\kappa_t\left(X_{t-1}+\varepsilon_t\right).$$

Thus its conjecture about ε_t is

$$\widetilde{\varepsilon}_t=\frac{i_t-i_{t-1}-\kappa_t X_{t-1}}{\kappa_t}.$$

Based on its observation of X_{t-1}, but prior to its observation of ε_t, the Fed decides on a value of k_t for that meeting. So its loss function is given by

$$E\left[\left(i_t^*-i_t\right)^2+\theta\widetilde{\varepsilon}_t^{\,2}\right]=\left(1-k_t\right)^2\left(X_{t-1}^2+\sigma_\varepsilon^2\right)+\theta\frac{1}{\kappa_t^2}\left(k_t^2\left(X_{t-1}^2+\sigma_\varepsilon^2\right)+\kappa_t^2X_{t-1}^2-2k_t\kappa_tX_{t-1}^2\right).$$

Differentiating with respect to k_t yields the first order condition

$$-\left(1-k_t\right)\left(X_{t-1}^2+\sigma_\varepsilon^2\right)+\theta\frac{1}{\kappa_t^2}\left(k_t\left(X_{t-1}^2+\sigma_\varepsilon^2\right)-\kappa_t X_{t-1}^2\right)=0.$$

In rational-expectations equilibrium, we have $\kappa = k$ so we have

$$k_t\left(1-k_t\right)=\theta\frac{\sigma_\varepsilon^2}{X_{t-1}^2+\sigma_\varepsilon^2}.$$

Future values of X_{t+n} are the accumulation of ε_{t+j} shocks, depreciated by the adjustment speeds k_{t+h}, so we have

$$X_{t+n}=\sum_{j=0}^{n}\left(\prod_{h=j}^{n}\left(1-k_{t+h}\right)\right)\varepsilon_{t+j}.$$

The future federal funds rate is the target rate i_{t+n}^* less X_{t+n}

$$\begin{aligned} i_{t+n} &= i_{t+n}^*-X_{t+n}=i_{t-1}+\sum_{j=0}^{n}\varepsilon_{t+j}-X_{t+n} \\ &= i_{t-1}+\sum_{j=0}^{n}\left(1-\left(\prod_{h=j}^{n}\left(1-k_{t+h}\right)\right)\right)\varepsilon_{t+j}. \end{aligned} \tag{A.7}$$

The variance of the future funds rate is

$$Var_t\left[i_{t+n}\right]=Var_t\left[\sum_{j=0}^{n}\left(1-\left(\prod_{h=j}^{n}\left(1-k_{t+h}\right)\right)\right)\varepsilon_{t+j}\right]$$

Note that the coefficient on ε_{t+j} is increasing in k_{t+h}. Thus increasing k_{t+h} increases the variance all else equal. However, given that $0<k_{t+h}<1$, we have

$$\prod_{h=j}^{n}\left(1-k_{t+h}\right)\to 0$$

as $n\to\infty$, so the effect of k_{t+h} on the variance dies out at long horizons.

A.6 Proof of Proposition 7

We begin by characterizing the behavior of an appointed central banker whose concern about the volatility of the α-pseudorate is parameterized by θ_c. Given θ_c, ε_t , and X_{t-1}, this central banker follows a rule of the form

$$i_t=i_{t-1}+k_t\left(\theta_c,X_{t-1}\right)\cdot\left(X_{t-1}+\varepsilon_t\right),$$

while the market conjectures that the central banker is following a rule of the form

$$i_t=i_{t-1}+\kappa_t\left(\theta_c,X_{t-1}\right)\cdot\left(X_{t-1}+\varepsilon_t\right).$$

Given this, the market's conjecture about ε_t is

$$\widetilde{\varepsilon}_t=\frac{\left(k_t-\kappa_t\right)X_{t-1}+k_t\varepsilon_t}{\kappa_t}$$

and it sets the infinite-horizon forward rate to

$$i_t^{\infty} = i_{t-1} + X_{t-1} + \widetilde{\varepsilon}_t.$$

Thus, the change in α-pseudorate is

$$\Delta i_t^{psu} = \alpha\left(i_t - i_{t-1}\right) + (1-\alpha)\,\widetilde{\varepsilon}_t.$$

The central banker picks $k_t\left(\theta_c, X_{t-1}\right)$ to minimize the loss function

$$\begin{aligned} E_t\left[\left(i_t^* - i_t\right)^2 + \theta_c\left(\Delta i_t^{psu}\right)^2\right] &= (1-k_t)^2\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \theta_c\alpha^2 k_t^2\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) \\ &\quad + \theta_c\frac{(1-\alpha)^2}{\kappa_t^2}\left(\left(k_t - \kappa_t\right)^2 X_{t-1}^2 + k_t^2\sigma_\varepsilon^2\right) \\ &\quad + 2\theta_c\alpha\,(1-\alpha)\,\frac{k_t}{\kappa_t}\left(k_t - \kappa_t\right)X_{t-1}^2 + 2\theta_c\alpha\,(1-\alpha)\,\frac{k_t^2}{\kappa_t}\sigma_\varepsilon^2. \end{aligned}$$

Differentiating with respect to k_t yields the first order condition

$$\begin{aligned} 0 &= -\left(1-k_t\right)\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \theta_c\alpha^2 k_t\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \theta_c\frac{(1-\alpha)^2}{\kappa_t^2}\left(\left(k_t - \kappa_t\right)X_{t-1}^2 + k_t\sigma_\varepsilon^2\right) \\ &\quad + \theta_c\alpha\,(1-\alpha)\,\frac{1}{\kappa_t}\left(2k_t - \kappa_t\right)X_{t-1}^2 + 2\theta_c\alpha\,(1-\alpha)\,\frac{k_t}{\kappa_t}\sigma_\varepsilon^2. \end{aligned}$$

In rational expectations we have $\kappa_t = k_t$, so this reduces to

$$0 = -\left(1-k_t\right)\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \underbrace{\theta_c\alpha^2 k_t\left(X_{t-1}^2 + \sigma_\varepsilon^2\right)}_{\text{A}} + \underbrace{\theta_c\frac{(1-\alpha)^2}{k_t}\sigma_\varepsilon^2}_{\text{B}} + \underbrace{\theta_c\alpha\,(1-\alpha)\,X_{t-1}^2}_{\text{C}} + \underbrace{2\theta_c\alpha\,(1-\alpha)\,\sigma_\varepsilon^2}_{\text{D}}. \tag{A.8}$$

Term A reflects the direct effect of changing k_t on the short-rate component of the pseudorate. Term B reflects the effect of changing k_t on the inferred shock $\widetilde{\varepsilon}_t$, and term C reflects the interaction of the two effects. Note that (A.8) implies that

$$\frac{\partial k_t}{\partial\theta_c} = \left[\begin{array}{c} -\left(\left(X_{t-1}^2 + \sigma_\varepsilon^2\right)\left(1 + \theta_c\alpha^2 k_t\right) - \theta_c\frac{(1-\alpha)^2}{k_t^2}\sigma_\varepsilon^2\right)^{-1} \\ \times\left(\alpha^2 k_t\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \frac{(1-\alpha)^2}{k_t}\sigma_\varepsilon^2 + \theta_c\alpha\,(1-\alpha)\,X_{t-1}^2 + 2\theta_c\alpha\,(1-\alpha)\,\sigma_\varepsilon^2\right) \end{array}\right] < 0$$

Society's ex ante loss is given by

$$\begin{aligned} E_t\left[\left(i_t^* - i_t\right)^2 + \theta\left(\Delta i_t^{psu}\right)^2\right] &= \left(1 - k_t\left(\theta_c\right)\right)^2\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \\ &\quad \theta\left(\alpha^2 k_t\left(\theta_c\right)^2\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + (1-\alpha)^2\sigma_\varepsilon^2 + 2\alpha\,(1-\alpha)\,k_t\left(\theta_c\right)\sigma_\varepsilon^2\right) \end{aligned}$$

Differentiating with respect to θ_c yields the first order condition

$$\left[-\left(1-k_t\right)\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \theta\alpha k_t\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \theta\alpha\,(1-\alpha)\,\sigma_\varepsilon^2\right]\frac{\partial k_t}{\partial\theta_c} = 0.$$

Evaluating at $\theta_c = \theta$ and using the appointed central banker's first order condition (A.8), we have that the term in the square brackets is

$$-\left(1-k_t\right)\left(X_{t-1}^2 + \sigma_\varepsilon^2\right) + \underbrace{\theta\alpha k_t\left(X_{t-1}^2 + \sigma_\varepsilon^2\right)}_{\text{A}} + \underbrace{\theta\alpha\,(1-\alpha)\,\sigma_\varepsilon^2}_{\text{D/2}} = \underbrace{-\theta\frac{(1-\alpha)^2}{k_t}\sigma_\varepsilon^2}_{\text{-B}} \underbrace{-\theta\alpha\,(1-\alpha)\left(X_{t-1}^2 + \sigma_\varepsilon^2\right)}_{\text{-(C+D/2)}} < 0.$$

Since $\partial k_t/\partial\theta_c < 0$, this implies that society's first order condition is positive at $\theta_c = \theta$ and its ex ante loss can be reduced by setting $\theta_c < \theta$.