October 13, 2015

Monetary Policy at the Lower Bound with Imperfect Information about r^*

David López-Salido, Christopher Gust, Benjamin K. Johannsen, and Robert Tetlow[1]

1. Introduction

Most central banks implement monetary policy by setting a target (or a range) for a nominal overnight interbank interest rate. But the level of the short-term policy rate does not, by itself, convey the stance of monetary policy, partly because spending decisions depend on real, rather than nominal, interest rates and partly because the level of the real rate needs to be compared to a benchmark level to determine if the stance of monetary policy is tight or loose. That benchmark, or r^*, is not observable and must therefore be estimated, and those estimates are subject to uncertainty stemming from statistical errors as well as model misspecification errors. As we discuss below, acknowledgement of this uncertainty is crucial for operationalizing the r^* concept in policy analysis, particularly when the policy rate is near its effective lower bound.

Many economists have argued that the recent financial crisis and Great Recession resulted in a decline in r^*, one that might be expected to persist for a long time. As discussed in the memo titled "Real Interest Rates over the Long-Run," shifts in low-frequency factors—such as long-run productivity growth, trend population and labor force participation, and in the trend capital-labor ratio—may affect longer-run r^*, with implications for monetary policy. In addition, a myriad of temporary but persistent "headwinds"—the effects of the financial crisis on credit conditions, the slow recovery in the housing sector, restrictive fiscal policies, and global financial strains—may themselves cause transitory movements in r^* away from its longer-run value, as discussed in the memo titled "Estimates of Short-Run r^* from DSGE Models." At the practical level, it may be difficult to disentangle factors inducing temporary fluctuations in r^* from those with longer-run effects. It follows that policymakers may misperceive the nature and persistence of changes in r^*.

This memo focuses on two questions. First, given that measurements of r^* are imperfect, how should policymakers use those measures when conducting monetary policy? And second, what

[1] In preparing this memo we have benefitted from the comments of Bill English, Thomas Laubach, Steve Meyer, Marc Giannoni, and Marco del Negro. Kathryn Holston has provided outstanding research assistance.

are the implications of uncertainty about r^* for using simple policy rules as guides for conducting monetary policy? Our analysis uses a simple New Keynesian (NK) model.[2] Our choice of model means that in terms of the taxonomy laid out in the memo on "r^*: Concepts, Measures and Uses," we are equating r^* with the level of the short-term *natural* real rate of interest.

We begin by studying how optimal discretionary policy (ODP) responds to a change in the natural rate of interest when policymakers have imperfect information about the magnitude of the change. Under discretion, policymakers cannot credibly commit to carrying out a plan that requires them to make future choices that would be suboptimal at future times. The discretion concept limits policymakers' ability to influence private-sector expectations regarding the federal funds rate and other variables. Instead, the private sector knows that future Committees will always re-optimize without regard for policymakers' past promises. We pay particular attention to the case where the effective lower bound (ELB) on nominal interest rates is initially binding, and we characterize the optimal path for the policy rate around the time of departure from the ELB. We next compare the macroeconomic outcomes under the ODP to outcomes when policy is governed by simple rules (Taylor (1999) and first-difference rules). For some of these comparisons, the coefficients of our simple rules are optimized, given our model economy, a step that ensures that the comparison with the ODP is evenhanded. We also discuss the performance of simple rules that have been modified to incorporate a time-varying intercept term as a proxy for changes in the natural real rate, as well as the communications challenges associated with committing to policy actions that follow a simple rule with a time-varying intercept term.

Broadly, the results highlight the interplay of the effective lower bound and incomplete information about the natural rate in shaping optimal discretion as a basis for pursuing a *risk-management approach* to policymaking. Under the risk management approach, policymakers make their policy rate decisions while considering the entire distribution of future outcomes, including the fact that certain choices in the future might not be feasible, because of the ELB, and thus choose policy settings today that reduce the cost of such outcomes. This approach to decision making can be thought of as "taking out insurance" against adverse outcomes.

[2] The main benefits of the NK model we employ are as follows: (1) it accounts for the nonlinearities associated with the effective lower bound; (2) it formalizes the relationship between several concepts of short-run r^*; (3) it can outline the monetary policy implications of imperfect knowledge about the natural rate of interest; and (4) it can account for implications of the stochastic shocks borne by the economy in the presence of uncertainty about economic conditions. The choice of model can be seen as a compromise between tractability and relevance in generating qualitative outcomes that may help guide important elements of the normalization of US monetary policy. Evans *et al*. (2015) use a similar model to study the effects of uncertainty about the natural rate on optimal discretionary policy near the effective lower bound. Our analysis is complementary, and extends their work by analyzing the information acquisition problem facing the central bank and also considering simple policy rules. The details of the model can be found in a companion technical appendix.

We summarize our principal findings as follows:

1. If policymakers have imperfect information about r^*, then there is an *attenuation* in ODP responses to perceived shifts in r^*. That is, when policymakers receive potentially erroneous information about changes in r^*, they optimally respond by less than would be the case given full information.
2. When the economy is close to, or already at, the ELB, policymakers optimally respond to incoming information about r^* *by even less* than they would during normal times. The reason for the relatively muted responses to signals about r^* is to "take out insurance" against situations in which their misperceptions about r^* might cause policy to be "too tight," the ELB makes the policymaker unable to offset those misperceptions in the next period, and discretionary policy makes the policymaker unable to offset those misperceptions by committing to easier policy in future periods.
3. The performance of the ODP at the ELB can be reasonably approximated by a Taylor-type rule with a constant intercept term if the feedback coefficient on inflation is set to a high value. Similarly, the performance of the ODP at the ELB can be approximated by a Taylor-type rule where the natural rate of interest is used as the intercept term and the feedback coefficients on output and inflation are kept at benchmark values. The use of a first-difference rule can also come close to performing as well as the ODP if the inflation coefficient is set to a high value.
4. Including unobservable variables, like a time-varying r^*, in simple rules introduces communication challenges for policymakers as they need to justify their policy actions by referring to information that cannot be verified by private agents.

2. Optimal Discretionary Policy with Imperfect Information about r^*

As was discussed in the memo on concepts of r^*, the NK literature has studied the monetary policy implications of the fundamental factors affecting the economy by translating those factors into a measure of the natural real rate and assessing how optimal monetary policy can be identified with the natural real rate; see, e.g., Woodford (2003) and Galí (2008). For the simplest linear NK model—a two-equation model with monopolistic competition, sticky prices, and shocks from a narrow range of sources—the optimal, natural, and efficient real rates are identical, in the absence of the ELB.[3] More complicated models in this class, such as models with sticky wages in addition to sticky prices and with shocks to the mark-up of prices over

[3] See the Gust *et al.* memo on "r^*: Concepts, Measures and Uses" for a discussion of the definitions of optimal, natural, and efficient real interest rates and the conditions under which they would all be the same.

costs, introduce trade-offs in monetary policy setting and therefore distinctions between these different concepts of r^*.[4]

The effective lower bound on short-term nominal interest rates introduces important nonlinearities that sharpen the distinction between the optimal and the natural real rates.[5] Near the ELB *any* shock—supply or demand—creates a tradeoff between inflation and output for the policymaker because of the restrictions on monetary policy actions implied by the ELB. Most of the extant literature considers the case in which policymakers have complete information about the nature of fluctuations in the natural real rate. But because the natural rate is unobservable, we consider uncertainty about the level of the natural real rate. Formally, we assume that policymakers observe only a noisy signal about the natural rate, so that in every period they have to filter incoming information about fundamental factors and at the same time set policy optimally. To simplify the analysis, we assume that policymakers learn the true value of r^* with a one-period lag. In all instances, policymakers are assumed to understand the structure of the model economy, including their own uncertainty.

We study optimal monetary policy under *discretion*, by which we mean the optimal plan for the short-term nominal interest rate that *does not require commitment* to future policy actions.[6] We could have alternatively studied optimal monetary policy under commitment, which allows policymakers to achieve better outcomes for current inflation and output by influencing future expectations. However, this implies strong assumptions, including that current policymakers can impose constraints on future policymakers.

Uncertainty about the level of the natural real rate interacts with the ELB in ways that have important policy implications. To demonstrate this fact, Figure 1 compares the ODP under incomplete information with that under complete information. The policymaker optimally sets the nominal policy rate (on the vertical axes) in response to a signal indicating that the natural rate has increased (the horizontal axes show the magnitude of that signal, measured in percentage points, relative to a baseline measure that varies). Where the two panels differ is in their assumptions regarding the initial value for the natural rate. In the upper panel, the economy is in "normal times," meaning that the (initial) level of the policy rate is well away from the ELB. In

[4] We will not explicitly refer to the *neutral* real rate because in our model the stance of monetary policy can be defined in terms of the deviations of the actual real rate from the *natural* real rate.

[5] See the original work by Eggertsson and Woodford (2003) and Levin *et al*. (2010) for further references.

[6] In the model, optimal policy under commitment will outperform optimal discretion since the former would imply that, through commitments on expected future short-term interest rates, the policy plan will affect long-term interest rates that are the key channel through which monetary policy strategies affect current inflation and output. Of course, the role of expectations implies that the benefits of commitment strategies rely on *credible* communication that allows the private sector to understand such commitments. In the absence of such understanding, it is likely that some of the benefits are weaker or could even be overturned: For example, commitment strategies yield no benefits, and would even be costly, if the commitments had no influence on long-term interest rates and aggregate demand; the benefits of such strategies would also be smaller if expectations adjusted only very slowly to the announced strategy.

the lower panel of the same figure (and several that follow) the policy rate is modestly constrained by the ELB; that is, even small changes in r^* will move the ODP away from the ELB. In both panels, we compare the central bank's optimal response to the signal of an upward shift in r^* when the policymaker has full information—so that the signal perfectly reveals the value of r^*—with the case of incomplete information, in which the signal is distorted by random noise. Regardless of initial conditions, the policymaker operating under incomplete information optimally *attenuates* the policy response under incomplete information (the red dot-dashed line), relative to the complete information case (the blue dashed line).[7,8]

The bottom panel considers the case where the natural rate is low enough to convince policymakers that the ELB is modestly binding (this scenario is labeled in the figure as the economy is "at the cusp of the ELB"). For a positive r^* signal of sufficient magnitude, the policy rate now departs from the ELB, although it can take a perceived innovation to r^* of some magnitude to bring about that result when policymakers make decisions based on noisy signals. The logic of this finding is straightforward: When the ELB is in play, in the sense that there is non-trivial likelihood of its becoming a binding constraint and thereby impairing the operation of conventional monetary policy, the optimal policy is to implement a strategy that limits the likelihood of returning to the ELB at a subsequent post-departure date, thus reducing the macroeconomic cost of a return, should it occur. Deferral of departure from the ELB is one viable step toward achieving these objectives.

In Figure 2 we isolate the implications of the ELB by comparing experiments where information is incomplete in all cases, but the ELB may or may not bind, depending on initial conditions.[9] Whereas Figure 1 compared complete and incomplete information, Figure 2 maintains the hypothesis that information is incomplete and examines how the effects of the ELB constraint interact with the level of rates. The top panel shows results for normal times, meaning conditions where the ELB is far enough away to be (probabilistically) ignored, and the output gap and inflation are close to their target values. In this case, given the calibration of the model, policymakers translate perceived or actual changes in r^* into the same change in the policy interest rate. The bottom panel repeats the exercise assuming that the inherited natural rate is low enough that the effective lower bound is just binding, initially. When policymakers account for the ELB (the red dot-dashed line) they react much less in response to positive signals about the natural rate than would have been the case in the absence of the ELB (the blue dashed line).

---

[7] Without making too much of the model calibration, under full information, the first-period response of the policy rate to a positive signal in r^* is one-for-one. The response under incomplete information has been calibrated to be one-fourth of the full-information response. These magnitudes are sensitive to model calibration, but the direction of the effect is not.

[8] The reaction function displayed in the figure shows different responses to the signal that under perfect information fully reveals the "true" shock to r^*. Without perfect information, only policymaker perceptions matter for determining the optimal policy, assuming that policymakers make use of all available information efficiently.

[9] It follows that the red dot-dashed line in Figure 2 is the same as the one in Figure 1.

In particular, the policymaker keeps the nominal rate at the lower bound unless the signal about the natural rate is large enough (about 0.7 percent in our exercise), whereas in the absence of the ELB, the policy rate would rise even for a signal of zero.[10] Of note is the fact that as the signal grows larger, the reaction of the policymaker who optimally takes into account the ELB approaches, nonlinearly, the unconstrained optimal policy. Taken together, Figures 1 and 2 show that incomplete information is sufficient to induce attenuation of policy, relative to the full-information case, and that the addition of concern about the ELB accentuates that tendency. In effect, uncertainty about r^* leads policymakers to keep the policy rate lower than would be the case under complete information. The reason for the relatively muted responses to signals about r^* is to "take out insurance" against situations in which their misperceptions about r^* might cause policy to be "too tight" and the ELB constrains the policymaker from offsetting those misperceptions in the next period.[11]

Figure 3 shows the policy implications of increased uncertainty about the natural rate by adding to Figure 2 an instance where uncertainty about r^* is greater.[12] The two panels contrast these new results with those reported in Figure 2 with the optimal policy response in the more uncertain economy (the solid black lines). As one might expect, the added uncertainty has only a marginal effect on policy responses to signals about r^* when the policy rate is far from the ELB, as shown in the upper panel. But the effect gets large as the ELB gets closer, as in the lower panel. The value of "insurance" rises with uncertainty about r^*. Staying at the ELB for longer when there is more uncertainty is akin to the risk-management approach described by Evans *et al.* (2015).

A complementary way of demonstrating the effects of uncertain natural rates near the ELB is to look at the distribution of economic outcomes. Figure 4 shows the distribution of the natural rate, the nominal policy rate, the output gap, and the inflation rate during normal times (the black solid lines) versus occasions when the ELB is just binding (the red dot-dashed lines).[13] As shown in the upper-left panel, in normal times the mean of this distribution is higher than when the economy is in the vicinity of the ELB. Not surprisingly then, during normal times the distribution of the nominal policy rate (the upper-right panel) is toward the right-hand side of the panel and itself looks normal. However, when economic conditions are such that the policy rate is initially at or near the ELB, there is a substantial probability that the nominal interest rate will

[10] The reason for this result is that ignoring the ELB shifts the expected value of inflation and output up because it removes the asymmetric downside risks associated with the ELB.
[11] This observation is in accord with the conclusions of Evans *et al*. (2015).
[12] More precisely, we increase the volatility of the underlying shocks (supply and demand) by 20 percent while keeping the same signal-to-noise ratio as in the previous exercise. Because uncertainty matters for optimal policy, this illustrates that such a policy is not certainty equivalent.
[13] For this figure, the density for the natural rate is constructed according to its exogenous (stochastic) distribution. The other densities are for outcomes at time t given r^* at time $t-1$.

remain at the ELB, as shown by the red vertical bar at zero in the upper-right panel.[14] The ELB constraint, and the skewed distribution for the policy rate it induces, imparts skewness on the output gap and inflation, as shown in the bottom panels. It is the adverse outcomes identified by the long negative tails of the latter two distributions that lead to the risk management approach to policy noted above; in the presence of incomplete information, this approach provides precautionary stimulus to offset the downward bias associated with the effective lower bound.

3. Simple Policy Rules with Incomplete Information about r^*

3.1. Why Simple Rules?

The analysis to this point has focused on ODP, meaning a strategy in which the policymaker sets the nominal rate optimally in response to incoming data on a period-by-period basis but does not commit to future actions in order to achieve better economic outcomes today by influencing private sector expectations of future monetary policy. Optimal policy under discretion in the neighborhood of the ELB prescribes a complex, nonlinear reaction function for the policy rate that may, in practice, be difficult to communicate and implement. An alternative approach might be to commit to a policy rule in which the policy rate responds to the state of the economy in a systematic manner. In this section, we turn to the potential benefits of committing to a rule-based approach in setting monetary policy within the same economic framework. That is, our analysis in this section builds around the simple NK model employed in Section 2 and, as before, takes into account the proximity of the ELB on nominal interest rates and the imperfect knowledge that the central bank faces regarding r^*.

A shift from ODP to policy based on a simple rule involves several considerations. A simple rule cannot, in general, be optimal because the rule will adjust policy in response to only a subset of the state variables that would be considered by the ODP. But a commitment to monetary policy governed by a simple rule can, in principle, affect private sector expectations in a helpful way that is amenable to policymakers' achieving their objectives. Whether the benefits of commitment outweigh the costs of suboptimal feedback is something that will vary from case to case. Even so, as Taylor and Williams (2011) emphasized, because simple rules adjust policy in response to a small list of variables that are central to entire classes of models, rather than optimizing with respect to the idiosyncratic features of one particular model, such rules may avoid some of the problems of model misspecification. Some simple rules restrict the variables to which policy responds only to those that are directly observable, trading off the sub-optimality of this restriction when measurement is reliable against the benefits of avoiding errors when

[14] The thick vertical bar in the upper-right panel represents the probability of the short-term nominal interest rate being at the effective lower bound, while the lines that plot the density are conditional on it being above the lower bound.

measurement turns out to be unreliable. In the context of this memo, this means that advocates of rules must take a stand on how to incorporate imperfect information about r^*.

In the next subsection, we build upon our results by first examining two versions of the Taylor (1999) rule. The assumption regarding r^* is the only difference between the two specifications. In one polar case, the central bank assumes that the constant intercept of the rule is the steady-state real rate; in the other polar case, the central bank acknowledges the short-term fluctuations in the natural rate and adjusts the intercept term in the Taylor rule, but takes into account its imperfect information about this unobserved variable by attenuating the adjustment as suggested by the earlier discussion of optimal policy in the face of uncertainty about r^*. Then in Subsection 3.3, we consider a Taylor rule with optimized coefficients.

Because of uncertainty surrounding r^*, some researchers and policymakers have advocated alternative rules that do not depend on measurement of r^*, such as the first-difference rule. In light of this, we devote Subsection 3.4 to an examination of a version of the first-difference rule studied by, among others, Orphanides and Williams (2002). In this case, to determine the change in the policy rate the policymaker relies on a now-cast of current inflation and output growth, about which the central bank has only imperfect information.

3.2. The Taylor rule

As we noted above, advocates of simple rules generally must decide how to incorporate information about r^*. Many rules, including Taylor (1993, 1999), simply assume that it is constant.

To be general, we specify the Taylor (1999) rule in the following way:

$$i_t = \max[0, E_t\{r_t^* + \alpha(r^* - r_t^*) + \pi_t + 0.5(\pi_t - 2) + (y_t - y_t^*)\}] \quad (1)$$

The notation E_t indicates that expectations are conditional on the information available to the policymaker at time t.[15] Two comments are in order. First, nested within our specification is the standard Taylor rule, which does not include a time-varying r_t^* and instead includes an intercept term equal to the long-run real rate (which was set in the original rule to 2 percent).[16] This corresponds to the case where $\alpha = 1$. Second, because the policymaker has incomplete information, expectations must be formed about the current-dated variables in the rule: inflation π_t, the output gap $(y_t - y_t^*)$, and the natural rate.

---

[15] At time t, the policymaker knows previous values of output, inflation and r^*, but does not see inflation and output in the current period, and only receives a noisy signal about r_t^*.

[16] Gust *et al.* memo on "r^*: Concepts, Measures and Uses" and the memo by Kei-Mu Yi and Jing Zhang, titled "Real Interest Rates over the Long-Run" offer discussions both on the definition and the current estimates of the long-run real rate.

The left-hand column of Figure 5 compares the Taylor (1999) rule with $\alpha = 1$, the black solid line, to the ODP when the economy is in the vicinity of the effective lower bound. As before, the upper panel displays the policymaker's reaction function to signals about r_t^*. Even though r_t^* does not appear directly in the specification of the rule—because the intercept, in this instance, corresponds to the steady-state real rate and not the natural rate of interest—policymakers react to signals about the underlying shocks affecting r_t^* through associated changes in the expected rate of inflation and the expected output gap. The prescribed response of the Taylor (1999) rule is familiar to readers of the recent versions of the MPS section of Tealbook B: the rule counsels notably higher nominal interest rates than the ODP *even when the signal is low*; indeed, the policy rate departs from the ELB for any signal of r_t^* shown. Accordingly, near the ELB, the systematically high short-term nominal interest rate prescribed by the Taylor (1999) rule with a constant intercept results in wide dispersions of the outcomes for the output gap and inflation (shown in the lower two panels of the left-hand column) and average outcomes that are well below target.[17]

It is arguably suboptimal for a policymaker to not react to changes in the natural real rate, especially if changes in the natural rate can be seen as driving movements in the output gap and inflation. To address that concern, the blue dashed lines in the same column of the figure show the outcomes when $\alpha = 0$, meaning that the intercept of the rule is allowed to change, one-for-one, with perceived movements in the natural rate. The upper panel shows that policymakers' reactions to signals about r_t^* are similar to the optimal discretionary policy (the red dot-dashed line). And because policymakers respond directly to signals about r_t^*, economic outcomes are notably closer to the ODP results. As such, the distributions of outcomes for the output gap and inflation resemble the distributions under discretion and are centered near their target levels (shown in the lower panels).

As comforting as this result might appear, it bears noting that introducing a time-varying intercept into the Taylor (1999) rule—particularly an intercept that entertains as much variability as does the natural real rate—is subject to the understandable criticism that the natural rate is unobservable and difficult to measure. It follows that even if policymakers commit to such a rule, its behavior could appear discretionary to the public, a fact that would substantially complicate communications and might impair public accountability.

[17] One might argue that one should not take as given the particular level of the long-run real rate of 2 percent from Taylor (1999), particularly in the context of a simple model like the one we use here. We have also tried the staff estimate for the long-run real rate of 1.5 percent; the results barely differed from those shown in Figure 5. Interestingly, a lower and constant r^* would also have the effects of making the economy more prone to return to the ELB, which would result in undesirable inflation and output outcomes.

3.3 An Optimized Taylor rule

An alternative to the regular Taylor (1999) rule, with its essentially *ad hoc* specification of the feedback coefficients on inflation and the output gap, is to optimize the rule coefficients, conditional on the model, the form of the rule including the presence of the ELB, and the shocks.[18] Optimizing the coefficients in this way ensures that whatever performance is turned in by the rule is not an artifact of an arbitrary calibration, which provides a solid foundation for comparing the performance of the simple rule with the ODP. As it happens, in this model, the optimized coefficient on inflation turns out to be quite large, while the coefficient on the output gap is little different from that in the regular Taylor (1999) rule. The green solid line in the right-hand column of Figure 5 shows the results for the optimized Taylor rule. These results can be compared with other outcomes, including the ODP, the red dot-dashed line. As can be seen, the optimized Taylor rule advises that the policy rate stay at the ELB for a wider range of signals about r^* than does the ODP, but with a steeper response when departure does occur. The performance rendered by the optimized Taylor rule, while inferior to the ODP, is much better than that of the regular Taylor rule with a constant intercept (the black solid line in the left-hand column). More generally, while not as efficacious as the ODP, the optimized Taylor rule performs fairly well; its performance is remarkably similar to that of the regular Taylor rule with a time-varying intercept (the blue dashed lines in the left-hand column). Evidently, optimization of Taylor rule coefficients and allowing for a time-varying intercept term are close substitutes in their ability to provide accommodation under proper circumstances. Given that both policies respond to the same set of model shocks, just in somewhat different ways, the similarity in their performances is perhaps not so surprising.

3.4 A first-difference rule

Still another alternative to the simple Taylor (1999) rule that does not depend on r_t^* , equation (1), is a first-difference (FD) rule, our version of which is specified with *changes* in the policy rate on the left-hand side, and current output *growth* on the right-hand side (as well as inflation):

$$i_t = \max\left[0, i_{t-1} + E_t\{\gamma_\pi(\pi_t - 2) + \gamma_y\,(y_t - y_{t-1})\}\right] \qquad (2)$$

[18] More precisely, the optimized coefficients are those that minimize a (discounted) quadratic loss function written in terms of the output gap and the deviation of inflation from target, given the form of the rule and the shocks. Because of the ELB, the density forecast for the model needs to be computed to carry out the maximization; this optimization is doable for a small model such as the one we use here, but quickly becomes infeasible as the size of the model grows. That this optimization is for unconditional welfare means that the specifics of initial conditions, such as whether or not the economy is already at, or near, the ELB, is not taken into account in carrying out the optimization, which in turn means that the rule is not generic to those specific initial conditions.

where the feedback parameters, γ_π and γ_y, are optimized.[19] By specifying the policy rate in first differences, a rule like equation (2) eschews the dependence of policy on the natural rate of interest, which Orphanides and Williams (2002) have emphasized can, in some circumstances, outperform a simple rule like the Taylor rule that depends on unobservable levels of the natural rate. This advantage notwithstanding, it is still the case that the policymaker does not have complete information about current inflation and the current output gap, and thus must rely on potentially inaccurate now-casts of inflation and output. And because this specification carries forward any changes in the nominal interest rate, any misperceptions about inflation or output growth will persist for a long time.

The magenta dotted lines in the right-hand column of Figure 5 show outcomes for policy governed by the optimized FD rule, which, like the optimized Taylor rule, carries a fairly substantial feedback coefficient on inflation. Here, as before, the simulations are conducted for initial values of r^* where the ELB is initially (just) binding. The upper panel shows that the policy rate responds to noisy signals about the natural rate of interest in a fashion quite similar to the ODP. In essence, the shocks that would be encompassed by fluctuations in r^*, which are responded to under the ODP, are instead captured by changes in current inflation and output growth that appear on the right-hand side of equation (2). The similarity in (first-period) responses of the FD rule and ODP notwithstanding, there are noteworthy differences in performance across the two policies, shown in the lower panels in this column. There is more dispersion in economic outcomes under the FD rule—and, in particular, a higher incidence of positive output gaps and inflation above target—than under the ODP. This is because while the ODP tailors the policy response to the specific initial conditions of the day—the fact that the ELB is just binding in this instance—the FD rule is optimized for economic conditions and shocks, on average. Thus, in a greater set of circumstances under the FD rule than under ODP, policymakers will find themselves having not tightened enough, or as quickly, as they would have preferred had they benefitted from complete information. It is because the initial misperceptions are propagated for as long as they are that the FD rule turns in a somewhat less attractive performance than the ODP.

4. Conclusions

We began this memo with the premise that the use of some benchmark real rate might be useful for the conduct of monetary policy, at least under some conditions. Our point of departure was to consider carefully the implications of the uncertainty in the measurement of such a benchmark real rate for policy, and to do so in a concrete model-based fashion. To this end, we employed a

[19] The policymaker chooses the values of the parameters to minimize the discounted sum of weighted squared deviations of inflation relative to a 2 percent inflation objective and output relative to potential output.

simple New Keynesian model, and narrowed down the concept of r^* under study to be equivalent to the *natural rate of interest*, to study the uncertainty question both as it relates to optimal discretionary policy and to policy as governed by (modified) simple rules.

We documented that the unobservable nature of the natural rate of interest introduces inescapable uncertainties, and that these uncertainties become particularly relevant near the effective lower bound. Further, this uncertainty justifies attenuation in the policy response relative to the case when policymakers have full information. We noted that this attenuation, which can be thought of as a manifestation of a policy of risk management that "takes out insurance" against possible adverse outcomes, is heightened by the proximity of the effective lower bound. Finally, we examined how policy might adapt to r^* uncertainty, finding that, in our framework, a policy that increased the response coefficient on inflation in the context of simple rules, or one that allowed for time variation in the intercept term of a Taylor-type rule, could come reasonably close to replicating the performance of optimal discretionary policy.

Such findings always come with caveats, of course. In the current instance, besides the usual disclaimers that come with any model-based analysis, there is the issue of the implementability of each of the policies we studied. The elegance of simple rules is usually said to be their simplicity, and associated with that, their time invariance and robustness. Using a simple time-invariant rule—and sticking, more-or-less, to it—gives private agents the opportunity to learn how the rules works and formulate their expectations with that sort of systematic behavior in mind. Discretionary policy, even though it is optimal on a period-by-period basis, is a time-varying policy that does not share this feature. Taylor rules with time-varying intercept terms are, in many ways, an attempt to slip in through the back door something close to the optimality of the optimal discretionary policy without giving up on the simplicity and commitment of a simple rule. Even so, this approach substitutes the relative clarity of simple rules for the opaque process of estimating the natural real rate, an unobservable variable that may not be verifiable by the private sector. Whether a simple rule with a time-varying intercept term is any easier to communicate and adhere to than the optimal discretionary policy remains an open question.

References

Gust, C., Johannsen, B. López-Salido, and R. Tetlow (2015), "r^*: Concepts, Measures, and Uses," memorandum to the FOMC, October 13, 2015.

Eggertsson, Gauti B., and Michael Woodford (2003), "The Zero Interest-Rate Bound and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 1, 139-211.

Evans, Charles, Fisher, J. Gourio F. and S. Krane (2015), "Risk Management for Monetary Policy Near the Zero Lower Bound," *Brookings Papers on Economic Activity*, March 19, 2015.

Galí, J. (2008), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, Princeton University Press.

Levin, A., David Lopez-Salido, E. Nelson, and T. Yun (2010), "Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound," *International Journal of Central Banking*, 6 (1), 143-189.

Orphanides, Athanasios and John C. Williams (2002), "Robust Monetary Policy Rules with Unknown Natural Rates," *Brookings Papers on Economic Activity*, 2:2002, 63-118.

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

Taylor, John B. (1999), *Monetary Policy Rules*, (editor), University of Chicago Press.

Taylor, J. and John C. Williams (2011), "Simple and Robust Rules for Monetary Policy," in Benjamin Friedman and Michael Woodford(ed.), *Handbook of Monetary Economics*, Volume 3B, North-Holland, 2011, 829-860.

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

Figure 1: Optimal Policy Setting Under Alternative Information Sets

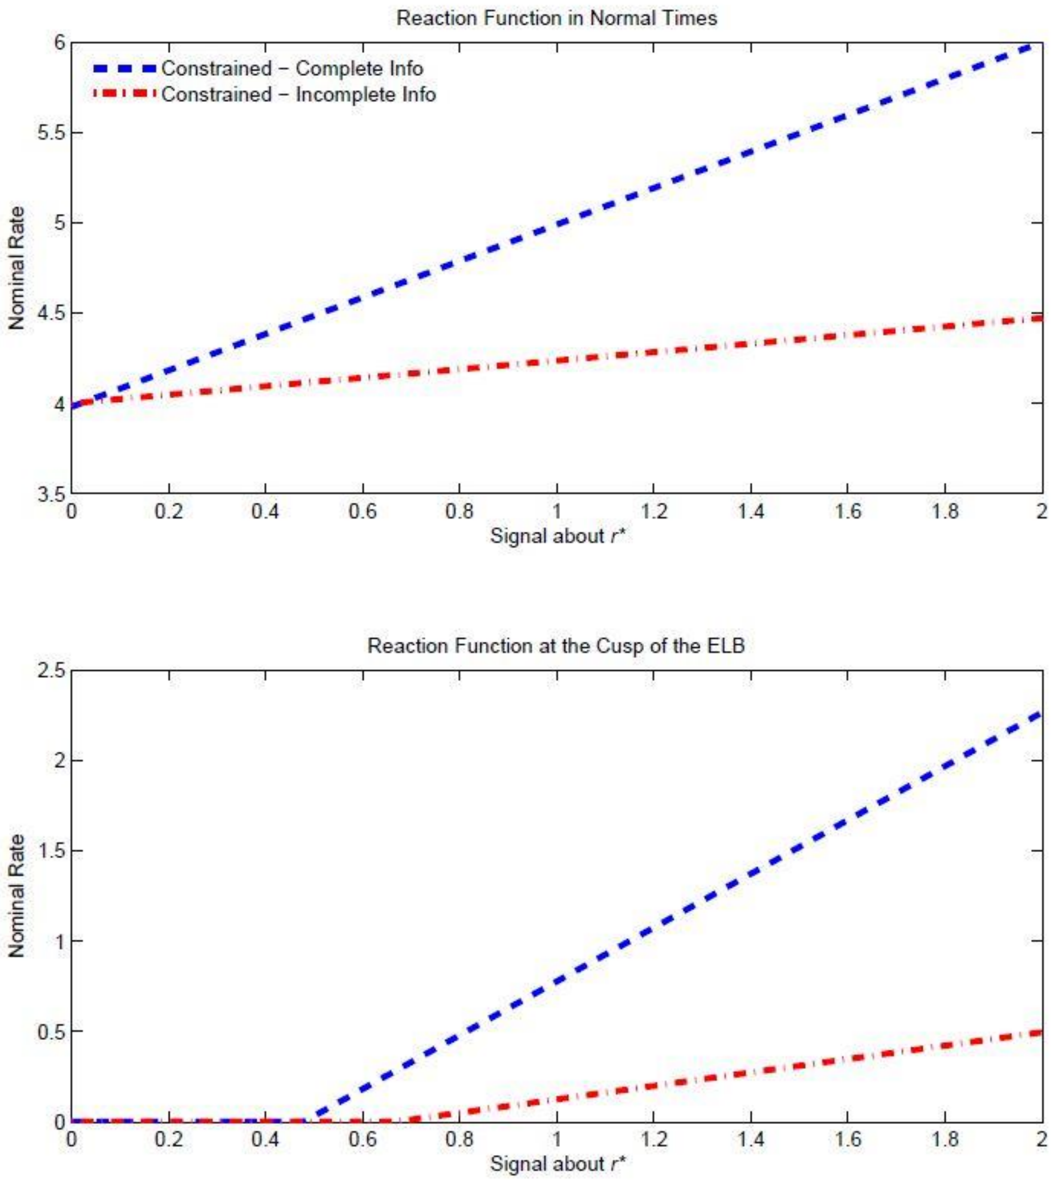

Figure 2: Optimal Policy Setting in the Proximity of the ELB

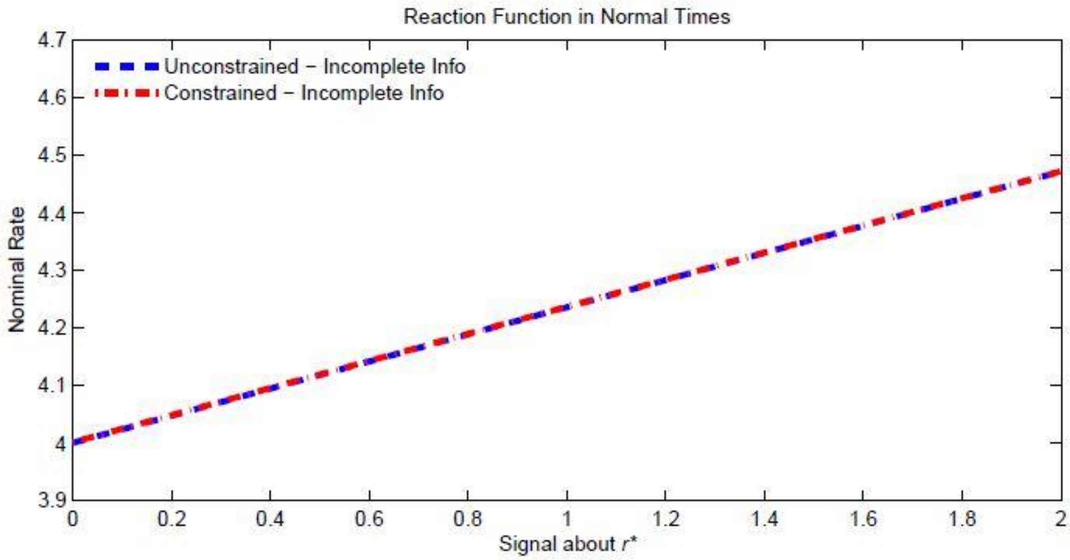

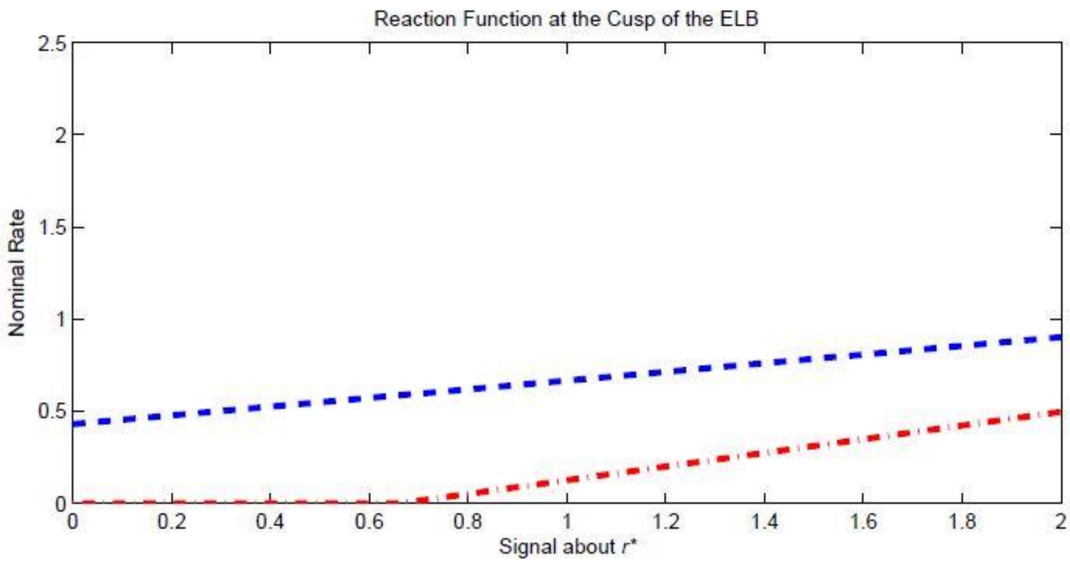

Figure 3: Optimal Policy Setting in the Proximity of the ELB
More Volatile Economy

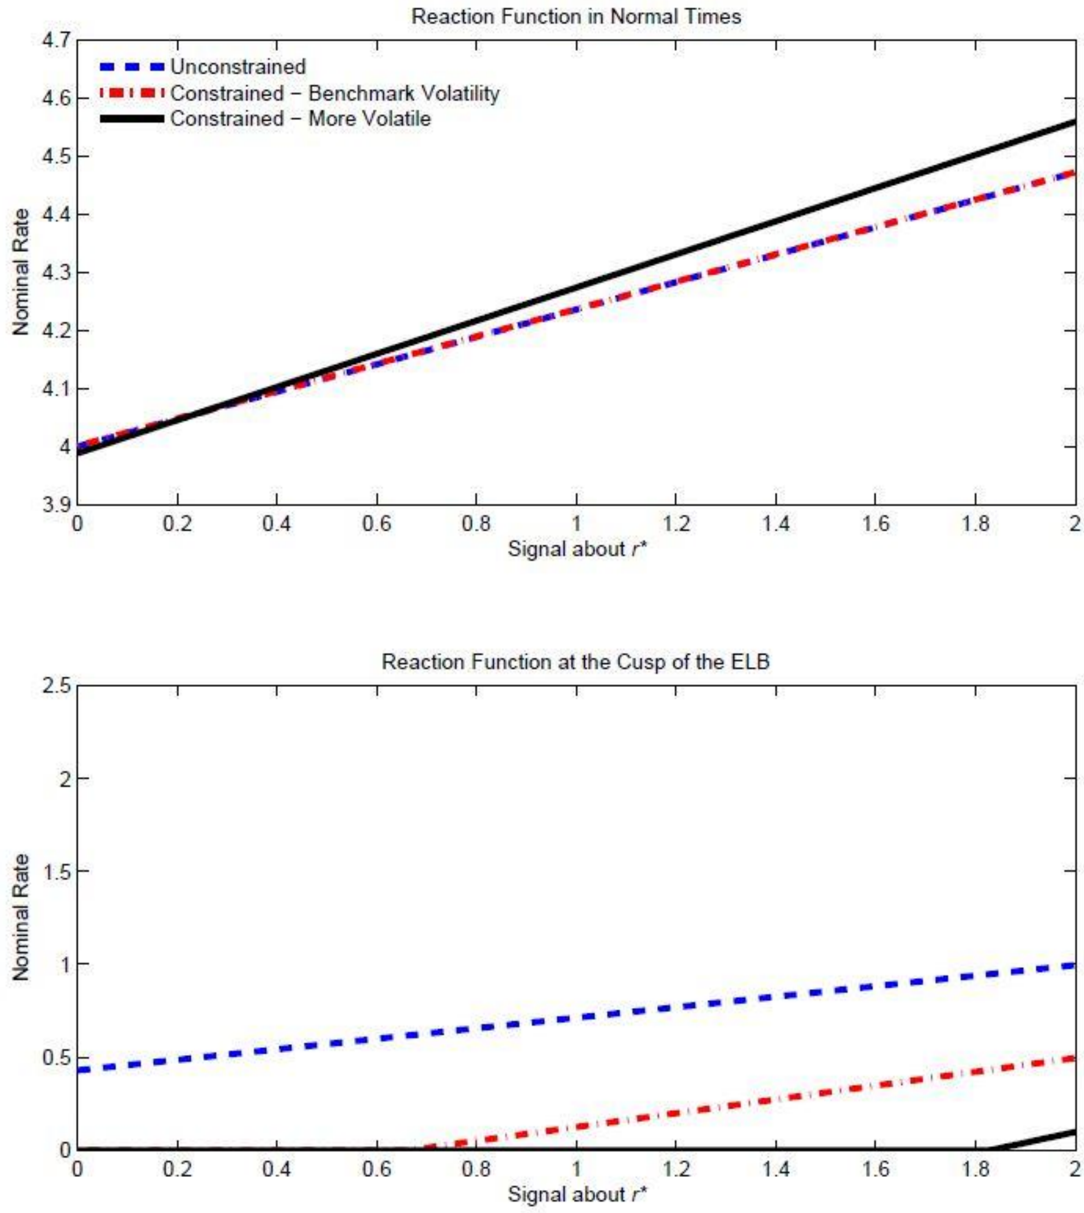

Figure 4: Distributions Conditional on Alternative Lagged Values of r^*

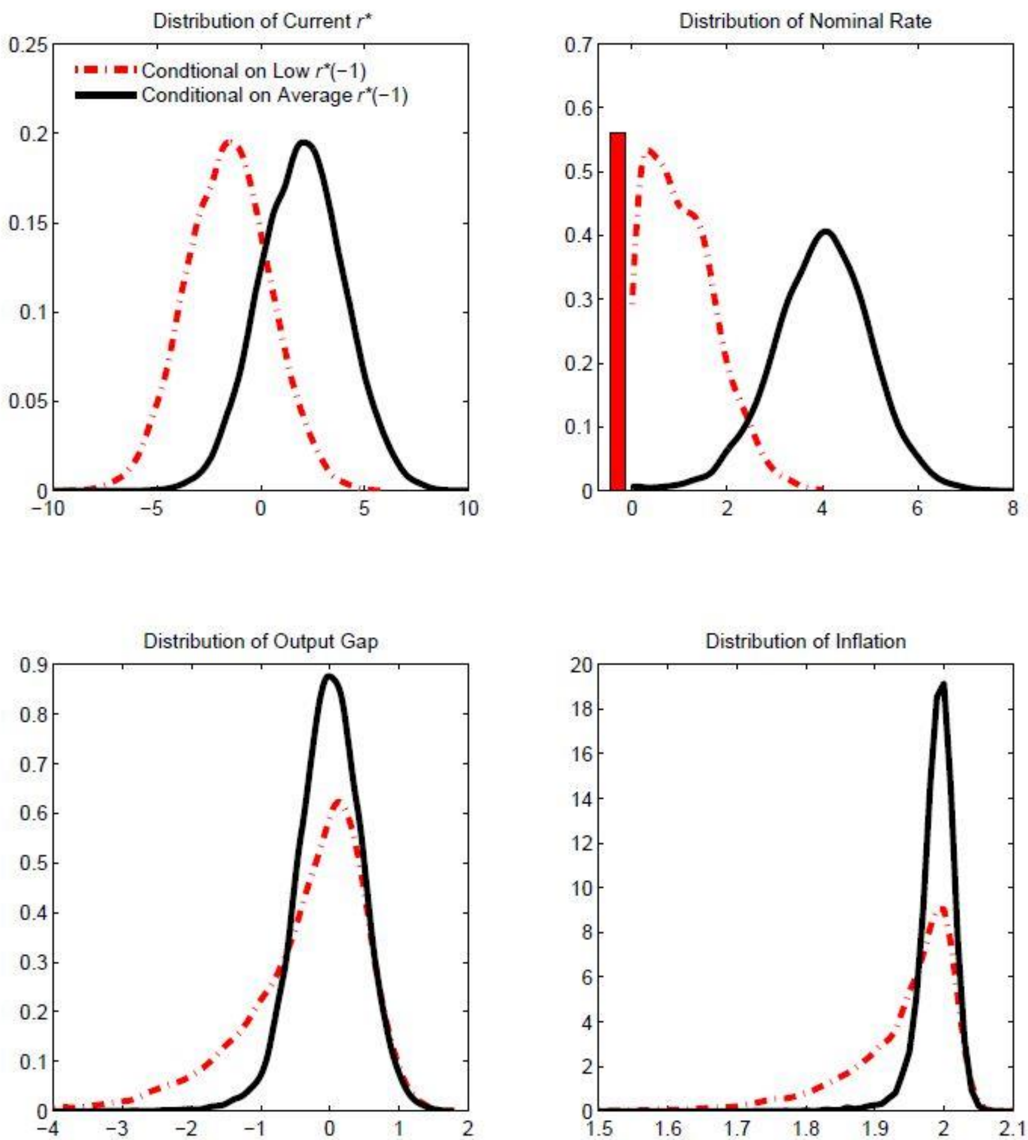

Figure 5: Comparison of Simple Rules and Optimal Discretion at the Cusp of the ELB