May 31, 2018

Martin Bodenstein
James Hebden
Fabian Winkler

Learning and Misperception: Implications for Monetary Policy Strategies

Introduction

A key determinant of the transmission of monetary policy to the economy is the public's expectation about the central bank's strategy for responding to economic developments. To assess the effects of alternative monetary policy strategies, Board staff routinely use macroeconomic models with specific assumptions about private-sector expectations formation. In the Monetary Policy Strategies (MPS) section of Tealbook A, it is assumed that private sector agents know the exact strategy pursued by the central bank and they believe that policymakers are committed to responding according to this strategy in all future circumstances. These assumptions may be reasonable when the monetary policy strategy is stable over time. However, these assumptions are more questionable when policy changes occur and the public does not immediately understand them.

In this memo, we expand the existing our conventional policy analysis to allow for the examination of changes in monetary policy strategies when the public needs to learn about those changes. We apply this learning framework to policy experiments like those routinely performed in the MPS. Our framework builds on an important academic literature, including Tetlow and von zur Muehlen (2001), and Cogley, Matthes and Sbordone (2015), which departs from rational expectations and assumes that the public infers key aspects of a policymaker's strategy using observed data.[1] More specifically, we assume that the public does not directly observe the parameters of the central bank's reaction function, and instead must form beliefs about likely current and future parameter values. To obtain estimates of these unknown parameters in the reaction function, the private sector agents run regressions using data on policy interest rates and their determinants (namely, measures of economic slack and price inflation). If agents learn the

[1] The literature on learning is extensive and distinguishes amongst others between procedural, rational, and educative learning. Our paper is closely related to Erceg and Levin (2003) since they also consider the situation in which the public is learning about aspects of the policy rule. In their model, the public is learning only about the inflation target in a rational manner. In contrast, we allow for the possibility that the public learns about other parameters of the rule, and our least-square learning mechanism falls into the class of procedural learning.

true nature of the policy change at a slow pace, their misperceptions about the course of monetary policy can lead to economic outcomes that differ notably from the case in which agents have instantaneous and complete information about the policy change.

Without loss of generality, we apply this framework to a specific policy change routinely analyzed in the MPS. In this case, the policy rule changes from the inertial version of the Taylor (1999) rule, which approximates the projected path of the nominal federal funds rate in the Tealbook baseline, to its non-inertial counterpart, also known as the balanced-approach Taylor (1999) rule. In recent years, the latter rule has prescribed a federal funds rate that remains considerably above the Tealbook baseline for several years, but this different policy path does not induce substantial differences in economic outcomes. This result is reasonable when the public understands the change in the central bank's policy reaction function (full information case) and believes that policymakers are committed to no longer engage in interest rate smoothing. In that case, the public correctly anticipates that the federal funds rate eventually will fall persistently below the prescriptions of the inertial Taylor (1999) rule. This anticipatory effect greatly mitigates the contractionary effects stemming from the large initial tightening associated with the balanced-approach Taylor (1999) rule.

However, when agents must learn about the change in monetary policy strategy from observed data (learning case), the initial unexpectedly fast tightening of monetary policy under the balanced-approach Taylor (1999) rule can lead to large contractionary effects. The public does not immediately attribute the observed higher level of the federal funds rate solely to the central bank's decision to no longer engage in interest rate smoothing. After all, the observed rise in the federal funds rate could also reflect the beginning of an aggressive tightening of policy or a one-time discretionary adjustment in the policy stance. If the public views some of the rise in the federal funds rate as signaling the beginning of a tightening cycle, then it can significantly push up the public's *anticipated* path of the federal funds rate and thus *realized* long-term interest rates. As a result, resource utilization and in particular inflation can experience a more pronounced fall than under the full information case routinely shown in the MPS.

Finally, our approach is general enough to be used in a broad class of linearized economic models. In fact, we confirm that the Small FRB/US (sfrb) model and the Estimated Dynamic Optimization-based (EDO) model deliver similar responses to changes in monetary policy strategy as the FRB/US model routinely used in the MPS. Moreover, the range of policy changes to which we can apply our framework is considerably broader than the experiments we consider in this memo. For instance, we can consider actual and perceived changes in the inflation target or the equilibrium real interest rate as well as the addition or replacement of economic variables in the policy reaction function (such as a switch to a rule targeting the price level instead of inflation). At the end of this memo, we discuss possible extensions and limitations to our framework.

Framework

We conduct our analysis in the Small FRB/US (sfrb) model and the Estimated Dynamic Optimization-based (EDO) model.[2] The monetary policy strategy of the central bank is given by an interest rate rule that prescribes a target value for the federal funds rate. Throughout our analysis, we distinguish between the *actual* policy rule—the interest rate rule from which the central bank derives its prescriptions—and the *perceived* policy rule—the interest rate rule that the private sector believes the central bank is following.

When expressed in deviations from the steady state of the model, the actual policy rule is of the form:

$$i_t = \rho_{i,t} i_{t-1} + (1 - \rho_{i,t})(\rho_{\pi,t}\pi_t + \rho_{y,t} ygap_t) + e_t.$$

The target value of the federal funds rate i_t is determined by the value of inflation relative to its long-run target value π_t, the output gap $ygap_t$, and the lagged value of the federal funds rate i_{t-1}. Our notation makes explicit that the parameters $\rho_t = \{\rho_{i,t}, \rho_{\pi,t}, \rho_{y,t}\}$ can in principle vary over time to accommodate changes in the central bank's systematic response to economic outcomes.

[2] Both models are linear and can be used to describe fluctuations of the economy around a baseline path. The sfrb model is closely related to the much larger and nonlinear FRB/US model. Like the FRB/US model employed in the MPS, the sfrb model features model-consistent expectation formation for inflation and asset prices under full information, but not for other variables. In the EDO model, by contrast, all expectations are model-consistent under full information.

The component e_t of the interest rate consists of a transitory and a persistent shock:

$$e_t = \varepsilon_t + \mathrm{v}_t$$

where ε_t is a white noise process, , and v_t follows an AR(1) process

$$\mathrm{v}_t = \rho_v \mathrm{v}_{t-1} + \eta_t$$

where the innovations η_t are again a white noise process. We see the transitory shock as representing a one-time discretionary adjustment to the federal funds rate, and the persistent shock as representing a long-lasting time-varying intercept term in the policy rule. Shifts in the public's beliefs about the intercept capture a variety of perceived persistent deviations from the policy rule, including those associated with the equilibrium real federal funds rate or the inflation target.[3]

Private sector agents know the structure of the economy as captured by the underlying model except for the details of the interest rate rule. They also observe all past and current realizations of the endogenous variables, in particular the values of the federal funds rate, inflation, and the output gap. While private sector agents know the general form of the interest rate rule, they do not necessarily know the vector of parameters ρ_t applied by the central bank. Neither do they observe past or current values of the two shocks, ε_t and v_t . Agents believe that shocks and changes to the parameters ρ_t are independent and normally distributed with variances

$$V(\varepsilon_t) = \sigma^2, \qquad V(\eta_t) = \tilde{\sigma}_{vt}^2\sigma^2, \qquad V(\rho_t - \rho_{t-1}) = \tilde{\Sigma}_t\sigma^2.$$

The private sector agents use their observations of economic data and statistical inference to inform their understanding of the central bank’s current and future monetary policy strategy, summarized in the perceived policy rule. The perceived policy rule resembles in form the actual policy rule with the parameters ρ_t being replaced by their sample estimates $\hat{\rho}_t$ and similarly for the transitory and persistent shocks. We denote the estimates of the two shocks by $\hat{\varepsilon}_t$ and $\hat{v}_t$, respectively.

[3] An increase in the value of the equilibrium real federal funds rate in the policy rule corresponds to an increase in the intercept, while an increase in the inflation target corresponds to a decrease in the intercept.

In a given period t, the public updates its beliefs regarding the properties of the rules as follows:

- Let the private sector's beliefs of the persistent shock and the rule parameters inherited from the previous period be denoted by $\hat{v}_{t-1}$ and $\hat{\rho}_{t-1}$.
- Within the period, differences arise between the observed value of the federal funds rate set by the policymaker under the actual interest rate rule and the value predicted by the perceived rule. Agents ascribe this difference partly to the transitory and partly to the persistent monetary policy shock. When computing the equilibrium prices and allocations in period t, we assume that agents form expectations about the future under anticipated utility (Kreps, 1998).[4]
- At the end of the period, agents update their beliefs about the parameters of the rule, the persistent shock, and the transitory shock in light of the newly observed data. To do so, they run a Bayesian regression.[5]
- The public's revisions to its past estimates, $\tilde{v}_t - \tilde{v}_{t-1}$ and $\hat{\rho}_t - \hat{\rho}_{t-1}$, depend on the public's beliefs about the volatility of changes in the persistent shock and the parameters relative to the transitory shock. These beliefs, which are embodied in the parameters $\tilde{\sigma}_{vt}^2$ and $\tilde{\Sigma}_t$, are an exogenous input to the learning process and can be time-varying.[6] If the public believes that large changes of a parameter are relatively likely in a given period, the Bayesian regression assigns a relatively large share of the discrepancy between the actual and the predicted federal funds rate in that period to a change in this specific parameter.
- In addition, the magnitude of the revisions to past estimates are larger when the discrepancy between the predicted and the actual value of the federal funds rate is larger.

A more detailed description of these steps can be found in a separate technical appendix to this memo.

---

[4] Anticipated utility refers to a widely used assumption in the learning literature that agents derive their decisions and their expectations about future economic developments under the assumption that their current perception of the policy rule parameters persists indefinitely. It is a simplifying assumption because at the same time, the public treats the parameters in the policy rule as random variables when learning about them. See Cogley and Sargent (2008) on interpreting anticipated utility as an approximation to Bayesian optimal learning.

[5] Assuming that the private sector agents update the parameters in the policy rule only at the end of the period simplifies the computational implementation, and is standard practice in the learning literature.

[6] The matrix Q_t can also be understood as the collection of signal-to-noise ratios in the Bayesian regression problem. Larger values of an element of Q_t indicate a higher signal-to-noise ratio for the corresponding parameter, and lead to relatively larger revisions of its estimate.

A question that arises naturally is whether the public's beliefs about the rule will eventually converge to the actual rule being followed by the central bank. For this to be the case, agents have to entertain the possibility of changes in those rule parameters that the central bank actually changes. But beyond that, it also has to be the case that the data agents observe are sufficient to identify the parameters that agents are uncertain about. In practice, this implies that we have to run stochastic simulations of the model, because deterministic simulations do not provide sufficient variation in the observables to identify the parameters of the policy rule.[7]

Application to the MPS

As noted above, we consider the economic effects of learning in a situation in which policymakers switch from a rule with inertia to a rule without inertia, using the sfrb model and the December 2017 Tealbook baseline. Up until 2017Q4, we assume that the central bank follows the inertial version of the Taylor (1999) rule and that agents previously learned the correct specification of the policy rule. In 2018Q1, the central bank switches from the inertial version of the Taylor (1999) rule to the balanced-approach Taylor (1999) rule by lowering the coefficient on the lagged value of the federal funds rate from 0.85 to zero. Expressed in terms of the notation used earlier, it is $\rho_{i,t} = 0.85$ for $t \leq$ 2017Q4 and $\rho_{i,t} = 0$ for $t \geq$ 2018Q1, while $\rho_{x,t} = \rho_x = 1$ and $\rho_{\pi,t} = \rho_\pi = 1.5$ for all t. To simplify the analysis, we further assume that agents know the true rule parameters with certainty up until 2017Q4, and also remain certain about the parameters ρ_π and ρ_x afterwards. Beginning in 2018Q1, they become uncertain and have to learn about the value of the inertia parameter $\rho_{i,t}$ and the intercept term.

To work through the effects of learning in this particular experiment, we start with Figure 1, which depicts how the private sector's beliefs about the values of the smoothing coefficient $\hat{\rho}_t^i$ (top left panel) and the persistent shock v_t (top right panel) evolve over time. At the beginning of 2018Q1, agents still believe that monetary policy follows the inertial version of the Taylor (1999) rule. However, by the end of that quarter, and after observing the realized values of the federal funds rate, inflation, and the output gap, agents attribute some of the unanticipated hike in

[7]Another condition for convergence is that the regression model that agents use has to produce unbiased estimates. Unbiasedness holds for the experiments described in this memo, but it would need to be verified case by case in different applications of our framework.

the federal funds rate to a lower value of the smoothing coefficient ρ_i. Initially, the updating step is small as agents view large changes to the smoothing coefficient as relatively unlikely. Given the uncertainty about the source of the abrupt federal funds rate increase, agents also raise their perceived value of the persistent shock. This misperception turns out to be crucial for the effects of learning because it translates into a persistently higher expected path of the federal funds rate. The misperception about the future path of the federal funds rate is illustrated in the bottom panel of Figure 1. The panel shows the paths of the federal funds rate as anticipated by the public in a given quarter, as well as the realized path of the federal funds rate. In 2018Q1, agents anticipate a path for the real federal funds rate that is considerably higher than the path that is eventually realized. The reason for the divergence between the anticipated and realized paths is precisely the misperception of the policy rule by the public, and in particular the public's estimate of the intercept in the rule.

Over time, agents adjust their beliefs about the smoothing parameter and the persistent shock towards their true values—that is, the values governing the actual policy rule implemented by the central bank. However, if the learning about the true value of the persistent shock is slow, misperceptions about the anticipated path of the federal funds rate will persist.

Our mechanism influences economic outcomes through its effects on the anticipated path of the federal funds rate. Figure 2 shows the evolution of the economy under three distinct assumptions about monetary policy: (1) Monetary policy follows the inertial version of the Taylor (1999) rule as in the Tealbook baseline (no switch, in blue). (2) Monetary policy switches to the balanced-approach Taylor (1999) rule and agents know the rule followed by the central bank (full information, in black). (3) Monetary policy switches to the balanced-approach Taylor (1999) rule, but agents have to learn about the change in the rule (learning, in red).[8]

The first two scenarios replicate the standard analysis shown in the MPS Section of the Tealbook. Under the Tealbook baseline (the blue line) the federal funds rate slowly rises from

[8] The lines shown in Figure 2 depict the average sample paths of the variables for stochastic simulations of the sfrb model around the Tealbook baseline. Unexpected deviations from the forecasted path of the economy facilitate private sector learning and guarantee that over time the agent's beliefs held about the monetary policy rule converge to the true rule.

1 1/4 percent towards 4 percent over the next 4 years. The sizable projected overshooting of output over its potential level occurs with inflation slowly returning to its long-run target of 2 percent. Whereas the balanced-approach Taylor (1999) rule calls for a significantly tighter path of the federal funds rate—the prescribed value of the federal funds rate immediately jumps above 3 percent—this monetary policy strategy induces a path for inflation that barely falls below the Tealbook baseline. At the same time, the positive output gap is only somewhat below its baseline path.

When the private sector agents do not understand the exact nature of the change in the monetary policy strategy that caused the rise in the nominal federal funds rate, the switch from the inertial version of the Taylor (1999) rule to the balanced-approach Taylor (1999) rule induces a larger reduction in inflation and real activity than under full information. With agents attributing some of the observed rise in the federal funds rate to a perceived increase in the persistent shock, the upward shift in the anticipated path of the federal funds rate shown in the bottom of Figure 1 induces a larger run-up in the real 10-year Treasury yield that persists for several quarters longer than under full information (shown in the bottom left panel of Figure 2). Inflation declines sharply and further increases in the output gap are forestalled. Over time, as agents review their current perception of the monetary policy strategy in the light of new data, they reduce their estimates of the intercept term and the smoothing term towards their true values. As this process occurs, the path for the 10 year yield converges to the path under full information and the paths for the output gap and inflation converge toward their full information paths as well.

The decline in inflation under learning is large relative to the fall in the output gap, which might seem surprising given that the Phillips curve in the sfrb (and FRB/US) model is relatively flat. To understand this result, it is important to remember that inflation in sfrb under model-consistent expectations is only to a small extent affected by the current output gap, and depends mostly on expected future inflation, which in turn depends on the anticipated output gap path in the future.[9] Figure 3 plots the anticipated paths of the real 10-year Treasury yield and the output gap under learning. In the quarters after the policy switch, agents anticipate long term real yields

[9] In the Risk and Uncertainty section of Tealbook, FRB/US is often simulated under VAR-based expectations for which these considerations do not apply.

to stay elevated, and therefore also expect the output gap to decline markedly for a prolonged time. These anticipated paths result from misperception about the persistent monetary shock and ultimately do not materialize. Nevertheless, they exert considerable downward pressure on the inflation rate in the short term, much more than the realized path of the output gap path reveals.

Discussion

The public's belief of the relative volatility of the parameters in the in policy rule and the policy shocks over time is a key determinant of the degree of misperception and the speed of learning in our model. The effects of a change in the monetary policy rule on the economy are stronger if agents attribute a larger share of the initial discrepancy between the perceived value of the federal funds rate and its actual value to the presence of a persistent shock. Figures 4 and 5 repeat our benchmark exercise when agents believe the (relative) volatility of the persistent shock to be larger than assumed previously. In the top right panel of Figure 4, the estimated persistent shock rises by three times as much as in Figure 1 and it remains at considerably higher values throughout the horizon shown. The anticipated path of the real federal funds rate suggests that agents expect even more tightening in the future and the real 10-year Treasury yield jumps up to 3.5 percent. Although the central bank cuts the federal funds rate almost back to its 2017Q4 level under its newly adopted balanced-approach Taylor (1999) rule, inflation drops down to 0.5 percent while the output gap narrows and the unemployment rate rises. As before, the decline in inflation is mostly due to anticipated declines in the output gap that are much larger than are ultimately realized.

Alternatively, if the public assigns zero probability to changes in the persistent shock, the change in the actual policy rule induces outcomes that closely resemble those under full information (not shown).

A key condition for our mechanism to matter quantitatively is that the change in the monetary policy strategy induces a sizable change in the observed value of the federal funds rate. If the inertial version of the Taylor (1999) rule and the balanced-approach Taylor (1999) rule prescribed similar values for the current value of the federal funds rate (for example, suppose the lagged value of the federal funds rate is close to the prescription of the balanced-approach Taylor

(1999) rule), then a change in the monetary policy rule followed by the central bank would only have a minor impact on the perceived intercept term and thus current and expected future long-term interest rates. Learning might still induce uncertainty about the true policy parameters, but this uncertainty would leave expectations of future policy rates little changed in response to the change in the policy strategy.

Our results are generally robust to the underlying economic model. In the EDO model, the switch of monetary policy strategy underlying Figures 1 and 2 described earlier induces a large drop in inflation under learning compared to the full information case. In comparison to the sfrb model, the EDO model features larger narrowing of the output gap and increase in the unemployment rate in response to the change in monetary policy strategy, as shown in Figure 6.

Extensions

Lowering the smoothing coefficient in the interest rate rule is only one possible way in which monetary policy strategies can change over time. Our approach can in principle be applied to all the simple rules currently shown in the MPS, including the price level targeting rule. Moreover, we have analyzed cases of private sector misperception of monetary policy after changes in the long-run inflation target (similar to Erceg and Levin, 2003) or in response a sequence of unanticipated discretionary adjustments in the federal funds rate. To save space, we refrain from including these experiments here.

One important caveat is that our current methodology only handles linear macroeconomic models. This implies in particular that our analysis abstracts from the zero lower bound on interest rates. Incorporating such nonlinearities would raise additional theoretical and numerical challenges. If all policy rules under consideration prescribed keeping the nominal federal funds rate at zero, agents would receive little or no signal from the policy rate about the change in monetary policy strategy. Switching to a policy rule that keeps the federal funds rate low for longer—is unlikely to provide as much easing of economic conditions under learning as under full information.[10]

[10] See, for example, Bodenstein, Hebden, and Nunes (2012) or Gust, Herbst, and Lopez-Salido (2018) who develop models of imperfect credibility at the effective lower bound.

Similarly, we have not applied our framework to the case of monetary policy strategies derived from optimal control settings. Learning in an optimal control environment introduces additional theoretical challenges, but offers the possibility of having the policymakers learn about unobserved changes to the structure of the economy.

References

Bodenstein, M., Hebden, J., and Nunes R., (2012). "Imperfect Credibility and the Zero Lower Bound," *Journal of Monetary Economics*, Elsevier, vol. 59, pages 135-149.

Cogley, T., Matthes, C., and Sbordone, A. M., (2015). "Optimized Taylor rules for disinflation when agents are learning," *Journal of Monetary Economics*, Elsevier, vol. 72(C), pages 131-147.

Cogley, T., and Sargent, T. J. (2008). "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making," *International Economic Review*, vol. 49(1), pages 185-221, February.

Erceg, C. J. and Levin, A. T., (2003). "Imperfect credibility and inflation persistence," *Journal of Monetary Economics*, Elsevier, vol. 50(4), pages 915-944, May.

Gust, C., Herbst, E., and Lopez-Salido, D., (2018). "Forward Guidance with Bayesian Learning and Estimation," mimeo.

Kreps, D., (1998). "Anticipated Utility and Dynamic Choice," 1997 Schwartz Lecture, in Frontiers of Research in Economic Theory, Edited by D.P. Jacobs, E. Kalai, and M. Kamien, Cambridge University Press, Cambridge, England.

Tetlow, R. J. and von zur Muehlen, P., (2001). "Robust monetary policy with misspecified models. Does model uncertainty always call for attenuated policy?," *Journal of Economic Dynamics and Control*, Elsevier, vol. 25(6-7), pages 911-949.

Figure 1: Policy expectations

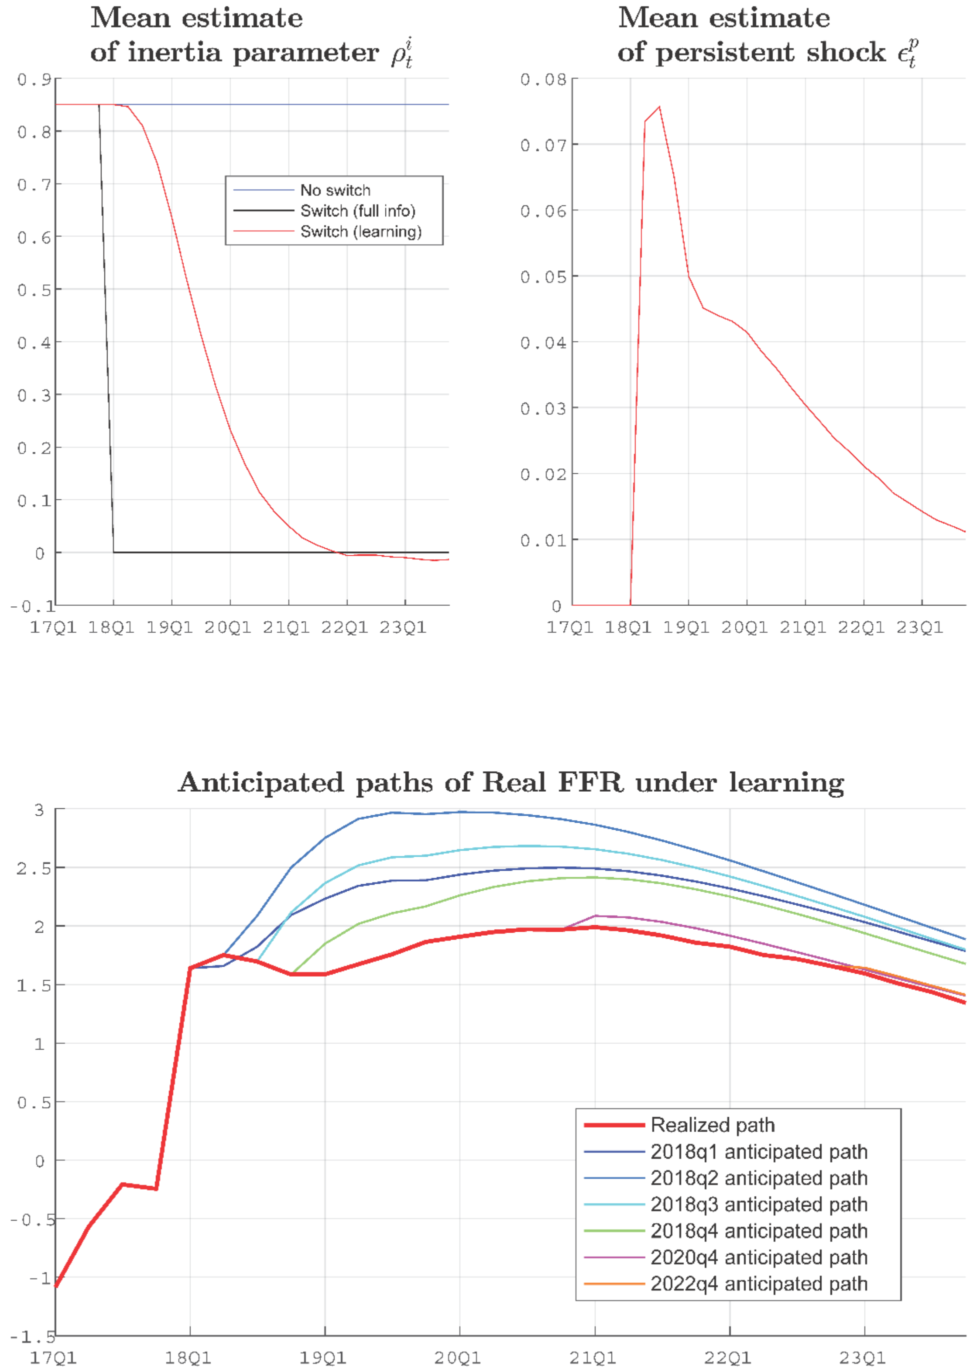

Figure 2: Outcomes

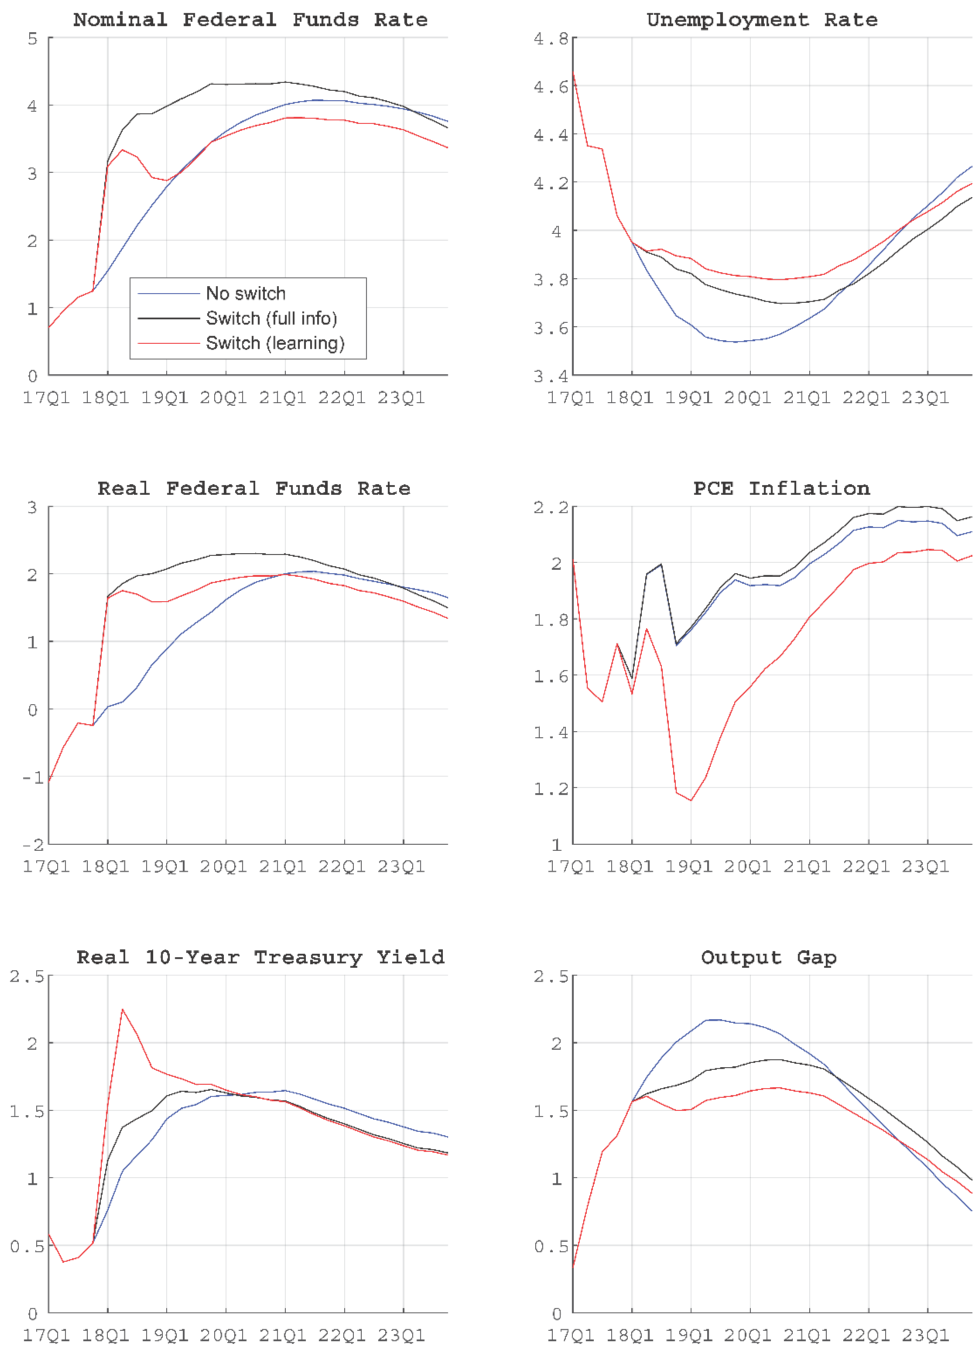

Figure 3: Long-term yield and output gap expectations

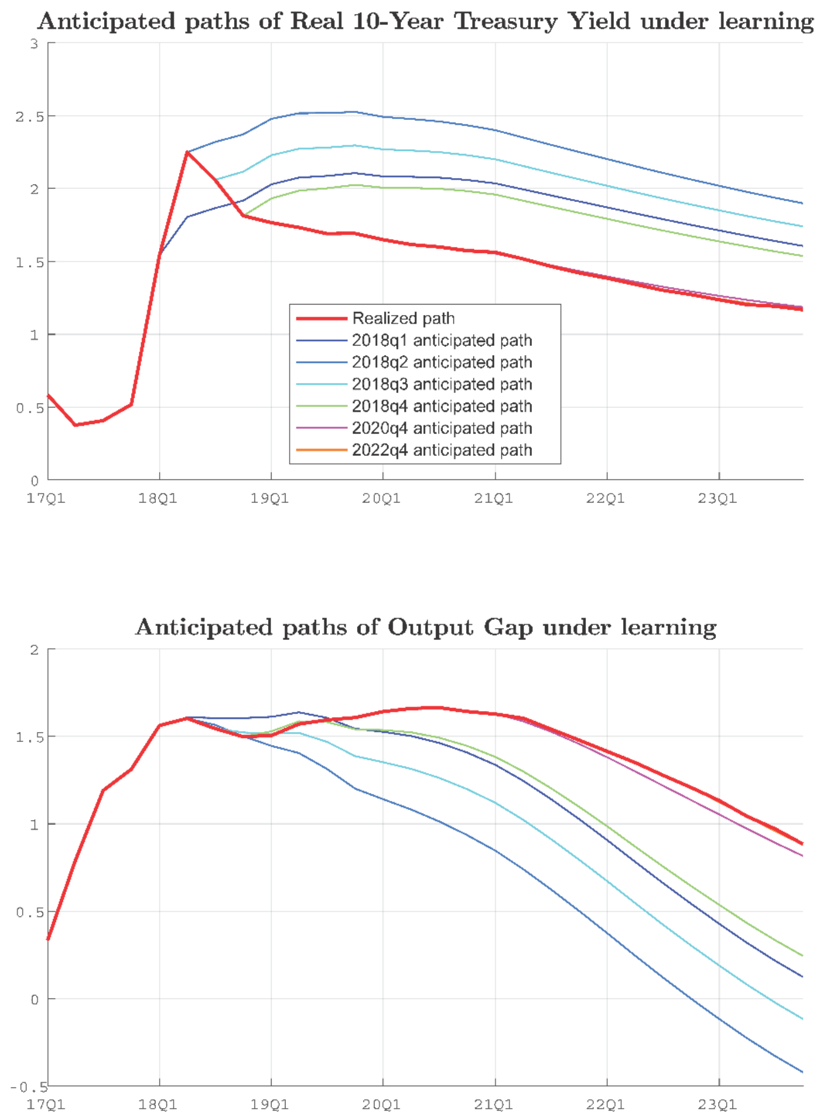

Figure 4: Policy expectations with increased misperception about the persistent shock

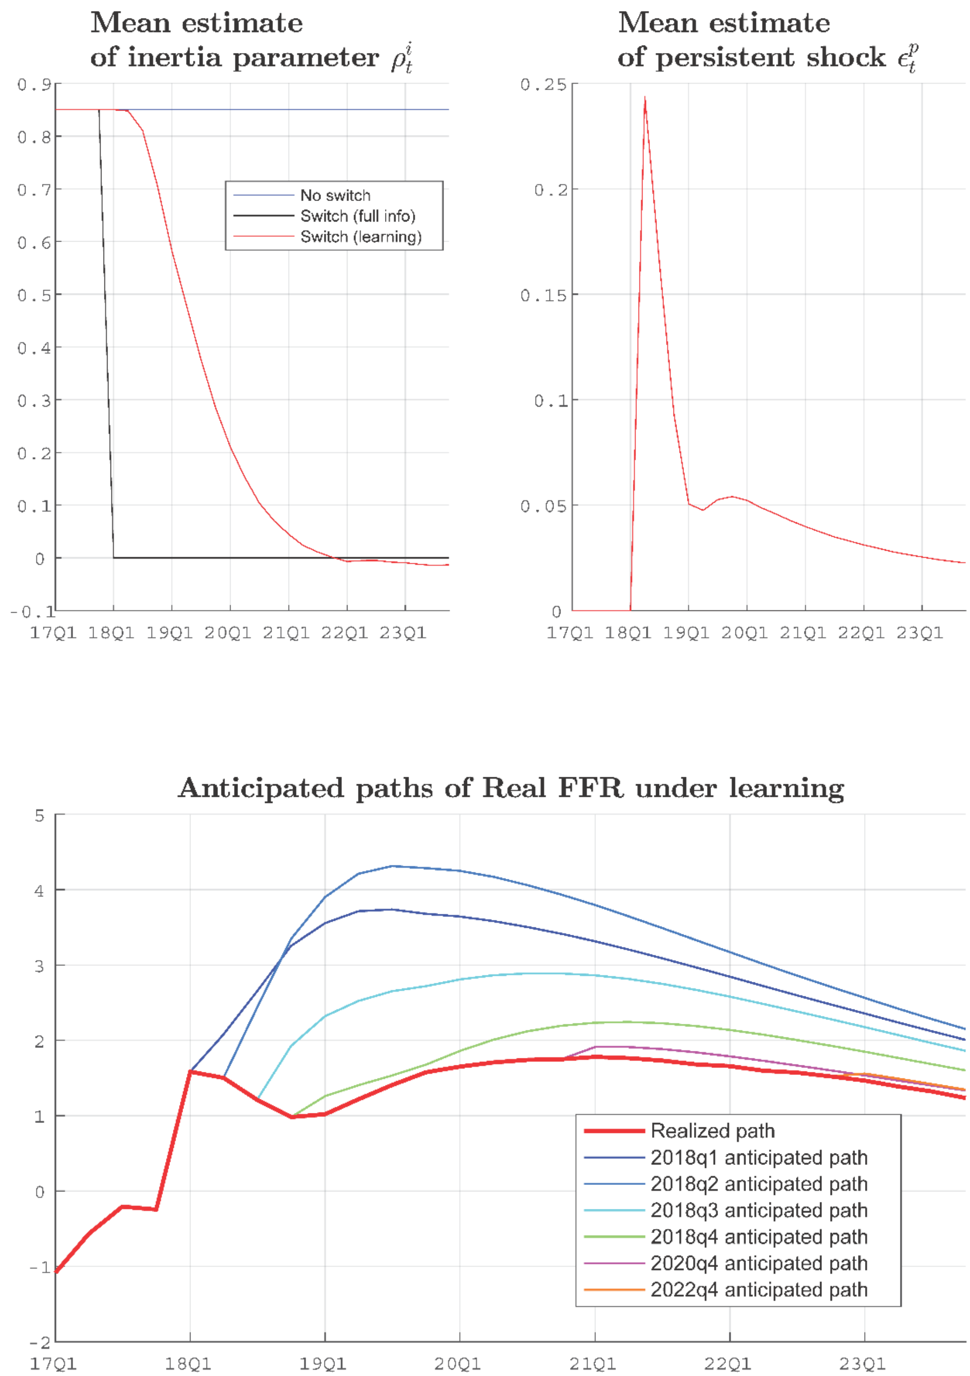

Figure 5: Outcomes with increased misperception about the persistent shock

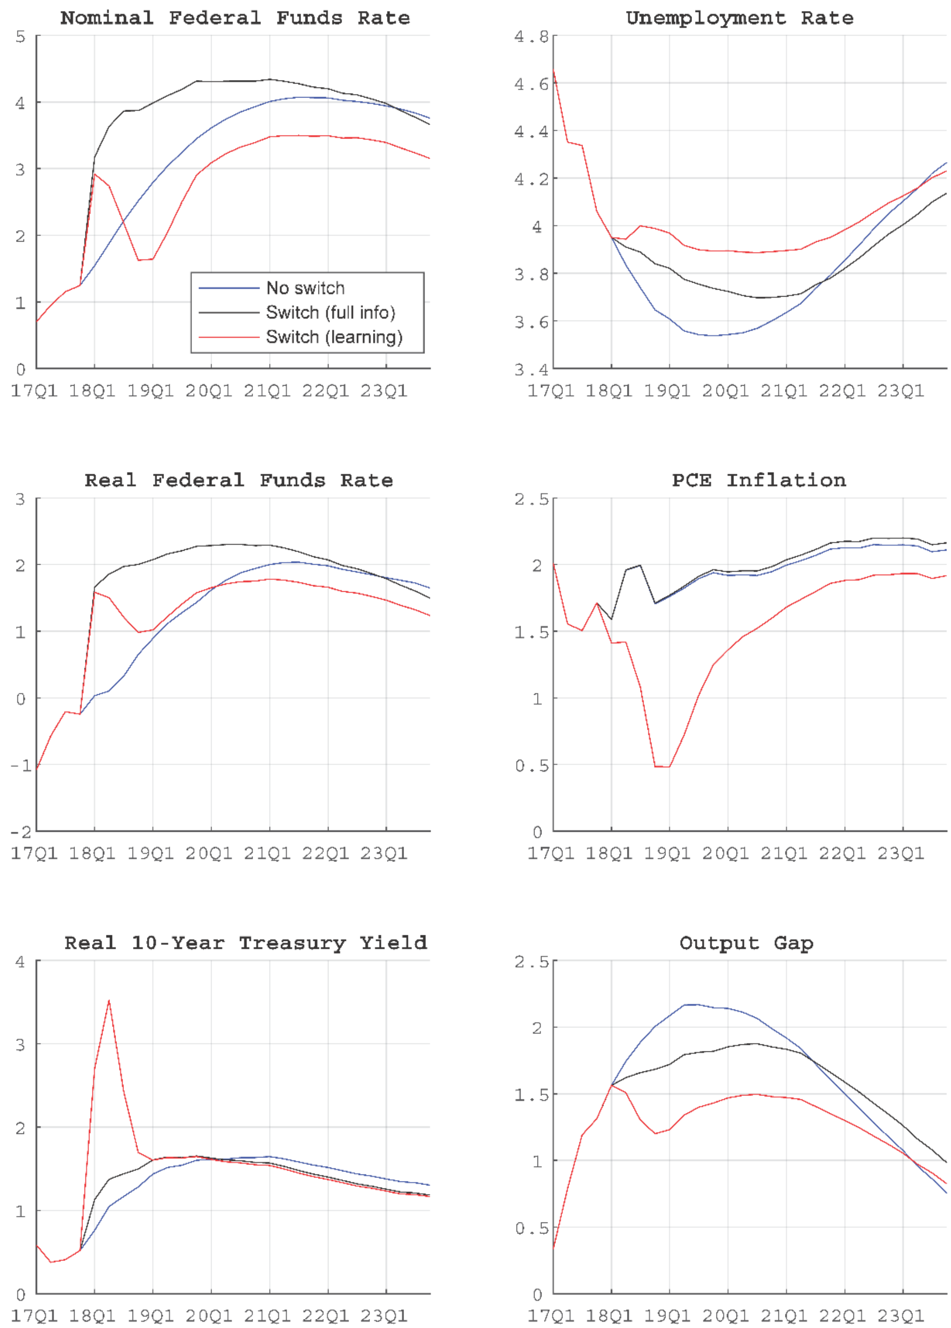

Figure 6: EDO model simulations

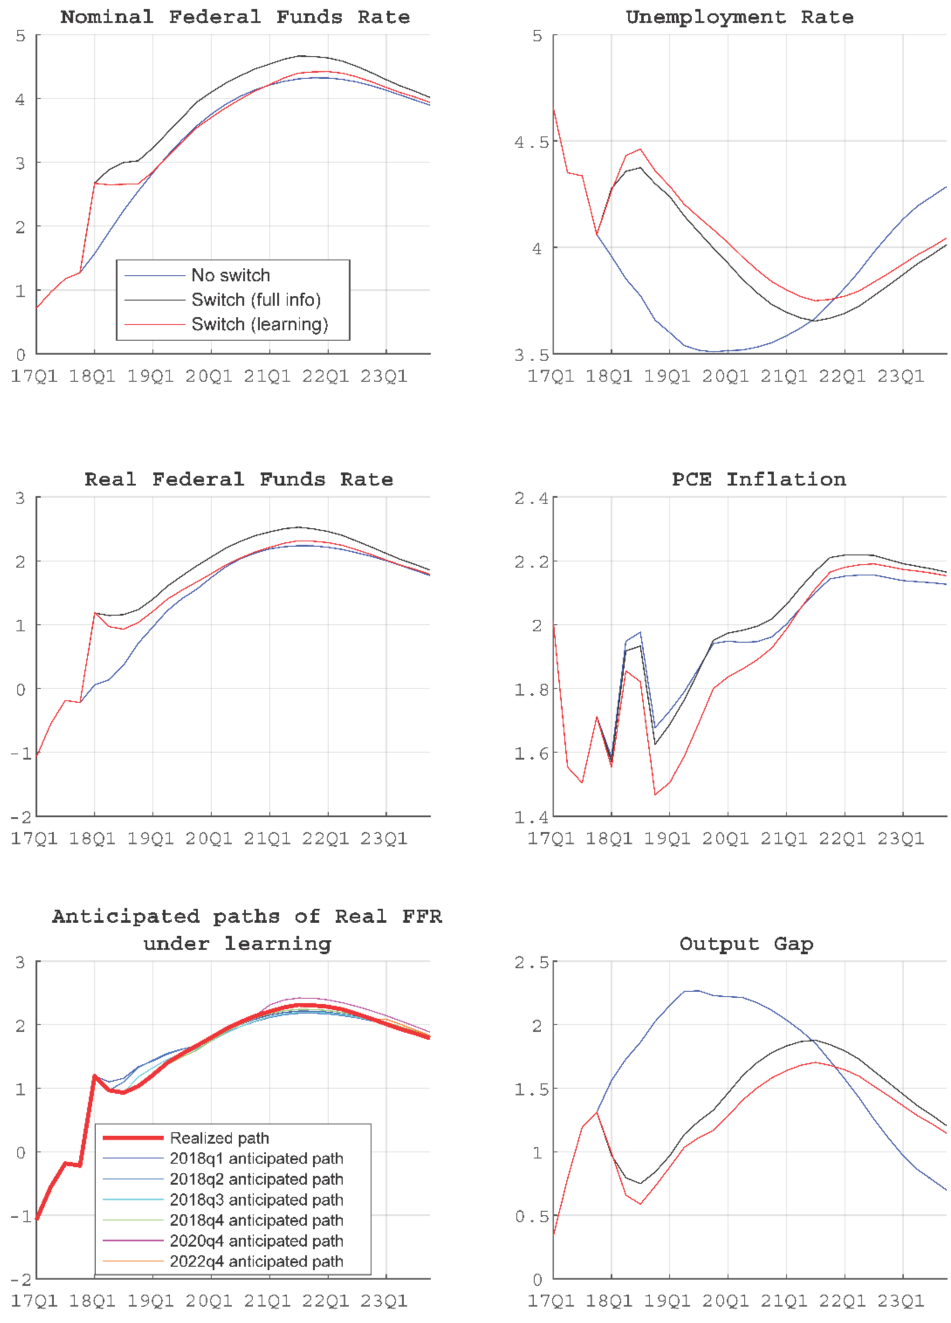

May 31, 2018

Martin Bodenstein
James Hebden
Fabian Winkler

Learning and Misperception: Implications for Monetary Policy Strategies – Technical Appendix –

1 Setup

We consider a standard linear rational expectations model with model variables y_t and with a central bank setting the nominal interest rate i_t according to an interest rate rule. The model equations are:

$$0 = F_1 \mathbb{E}_t y_{t+1} + F_0 y_t + F_{-1} y_{t-1} + F_i i_t + F_u u_t \tag{1}$$

$$i_t = \phi(\beta_t) x_t + v_t + \varepsilon_t \tag{2}$$

$$v_t = \rho_v v_{t-1} + \eta_t. \tag{3}$$

Equation 1 describes, through the matrices F, the set of all equilibrium conditions of the model except for the interest rate rule followed by the central bank. This interest rate rule is described by equation (2). It is linear in a subset x_t of the endogenous model variables y_t, and depends on potentially time-varying parameters β_t. It also contains a transitory policy shock ε_t and a persistent policy shock v_t described in equation 3 with innovation η_t. The other exogenous shocks in the economy are denoted u_t. All variables are expressed in deviations from a reference level, which is usually the steady state of the model.

Agents in the private sector only observe y_t, u_t and the nominal interest rate i_t, but they do not observe the monetary policy shocks. They also do not know some or all of the time-varying coefficients β_t and instead need to form subjective beliefs about some or all of the parameters in the interest rate rule. The operator $\mathbb{E}_t$ stands for rational expectations given information available to agents in time t.

In the application of the memo, the policy rule is an inertial Taylor rule of the form

$$i_t = \rho_{it} i_{t-1} + (1 - \rho_{it})(\rho_\pi \pi_t + \rho_y ygap_t) + v_t + \varepsilon_t$$

where all variables are expressed in deviation from their long-run levels. We assume that all parameters are known to the agents except for the interest rate smoothing parameter ρ_i. Then $\beta_t = \rho_{it}$, $x_t = (i_{t-1}, \pi_t, ygap_t)'$, $\phi(\beta_t) = (\beta_t, (1-\beta_t)\rho_\pi, (1-\beta_t)\rho_y)'$, and therefore

$$\phi(\beta_t) x_t = \beta_t (i_{t-1} - \rho_\pi \pi_t - \rho_y ygap_t) + \rho_\pi \pi_t + \rho_y ygap_t.$$

.

2 Belief Updating

At the end of each period, agents update their beliefs about the rule parameters ϕ_t and the policy shock v_t, using their observations of the endogenous variables and the interest rate. They do this using a Bayesian regression that can be cast as a Kalman filtering problem. The observation and state transition equations of the filter are:

$$i_t = \phi(\beta_t) x_t + v_t + \varepsilon_t \tag{4}$$
$$\beta_t = \beta_{t-1} + \epsilon_{\beta t} \tag{5}$$
$$v_t = \rho_v v_{t-1} + \eta_t. \tag{6}$$

We assume that agents believe that the transitory shock, innovations to the persistent shock, and changes to the parameters at time t are normally distributed white noise with:

$$\begin{pmatrix} \varepsilon_t \\ \eta_t \\ \beta_t - \beta_{t-1} \end{pmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} 1 & 0 & 0 \\ 0 & \tilde{\sigma}_{vt}^2 & 0 \\ 0 & 0 & \tilde{\Sigma}_t \end{bmatrix} \sigma^2\right). \tag{7}$$

We use the extended Kalman filter to describe the Bayesian filtering problem with the following updating equations:

$$\begin{pmatrix} \hat{v}_t \\ \hat{\beta}_t \end{pmatrix} = \begin{pmatrix} \rho_v \hat{v}_{t-1} \\ \hat{\beta}_{t-1} \end{pmatrix} + P_{t-1} \begin{pmatrix} 1 \\ h_t' \end{pmatrix} \left(i_t - \phi\left(\hat{\beta}_{t-1}\right) x_t - \rho_v \hat{v}_{t-1}\right) f_t^{-1} \tag{8}$$
$$P_t = \begin{pmatrix} \rho_v & 0 \\ 0 & I \end{pmatrix} \left(P_{t-1} - P_{t-1} \begin{pmatrix} 1 & h_t \\ h_t' & h_t' h_t \end{pmatrix} P_{t-1} f_t^{-1}\right) \begin{pmatrix} \rho_v & 0 \\ 0 & I \end{pmatrix} + \begin{pmatrix} \tilde{\sigma}_{vt}^2 & 0 \\ 0 & \tilde{\Sigma}_t \end{pmatrix} \tag{9}$$
$$f_t = \begin{pmatrix} 1 & h_t \end{pmatrix} P_{t-1} \begin{pmatrix} 1 \\ h_t' \end{pmatrix} + 1 \tag{10}$$
$$h_t = \frac{d\phi}{d\beta}\left(\hat{\beta}_{t-1}\right) x_t. \tag{11}$$

In our application in the memo, $h_t = i_{t-1} - \phi_\pi \pi_t - \phi_y ygap_t$.

The parameters $\tilde{\sigma}_{vt}^2$ and $\tilde{\Sigma}_t$ act as signal-to-noise ratios in the Bayesian updating equations. A larger value for $\tilde{\sigma}_{vt}^2$ indicates that forecast errors in the interest rate are more informative about $\hat{v}_t$, which will therefore get updated by a larger amount for a given forecast error of the interest rate. The updating gains also depend on the initial degree of parameter uncertainty as captured in the prior variance matrix P_0. In the first exercise of the memo, we set $\tilde{\sigma}_{vt}^2 = 0.01$, $\tilde{\Sigma}_t = 0.5$ for all $t \geq 1$, and P_0 as the 2x2-diagonal matrix with diagonal entries $P_{0,11} = 0.1$ and $P_{0,22} = 0.01$. These choices imply that agents believe that relative to the value of the smoothing coefficient, the value of the intercept is more uncertain at the start of the simulation but more stable over time. In the second exercise of the memo, the prior variance matrix is changed to $P_{0,11} = 0.5$, implying that agents are even more uncertain about the value of the intercept, leading to larger misperceptions of its value.

When $\rho_v = 0$, the persistent shock is in fact iid and not a latent state anymore. If additionally $\phi(\beta_t) y_t =$

$\beta_t' x_t$, then the updating equations simplify to:

$$\hat{\beta}_t = \hat{\beta}_{t-1} + P_{t-1} x_t \left(i_t - \hat{\beta}_{t-1}' x_t \right) f_t^{-1}$$
$$f_t = x_t' P_{t-1} x_t + 1$$
$$P_t = (P_{t-1} - P_{t-1} x_t x_t' P_{t-1}) + \tilde{\Sigma}_t$$

which is the Bayesian equivalent of the discounted recursive least squares (DRLS) learning in Tetlow and von zur Muehlen (2001). The variance matrix $\tilde{\Sigma}_t$ replaces the "forgetting factor" in the DRLS algorithm. For $\mid \tilde{\Sigma}_t \mid > 0$, the coefficients β_t are preceived to be time-varying, and the estimator $\hat{\beta}_t$ will overweigh the information from recent forecast errors and asymptotically discard information from the distant past. The prior variance matrix P_0 and the variance $\tilde{\Sigma}_t$ determine the size of the Kalman gain, i.e. the speed of belief updating. When $\tilde{\Sigma}_t = 0$, then the coefficients are perceived to be unchanged between $t-1$ and t.

Note that the regressors h_t have to be uncorrelated with the innovations ε_t and η_t for the Kalman filter to be a valid Bayesian estimator. Violation of this condition can lead to biased estimates. This condition is not innocuous because monetary policy shocks can easily be correlated with the variables that enter the policy rule. This is a well-known problem that arises when estimating reduced-form policy rules (cf. Clarida, Gali and Gertler, 1999).

3 Computing the equilibrium

We have described how agents update their beliefs at the end of each period, after observing equilibrium outcomes. We now describe how we compute these outcomes.

Each period, agents start with a belief about the rule coefficients $\hat{\phi}_{t-1} = \phi\left(\hat{\beta}_{t-1}\right)$ and the past value of the persistent policy shock $\hat{v}_{t-1}$. They then form expectations about other variables using the anticipated utility approximation (Kreps, 1998): Expectations about future variables in equation (1) are formed as if the rule coefficients were known with certainty and fixed forever at $\hat{\phi}_{t-1}$. The equilibrium also depends on agents' beliefs about the time-t values of the two policy shocks. These will be a function of the discrepancy between the observed interest rate and the prescription of the perceived policy rule. How much this discrepancy gets ascribed to the persistent shock v_t relative to the transitory shock ε_t depends on agents' beliefs about the signal-to-noise ratio in their signal extraction problem, which we describe here in detail.

Formally, in period t we solve the linear model under the assumption that $\phi_{t+s} = \hat{\phi}_{t-1} \, \forall s \geq 0$. The solution takes the form

$$y_t = A_t y_{t-1} + B_t u_t + c_{vt} \hat{\hat{v}}_t + c_{\varepsilon t} \hat{\hat{\varepsilon}}_t. \tag{12}$$

where $\hat{\hat{v}}_t$ and $\hat{\hat{\varepsilon}}_t$ denote agents' perceived values of the two monetary shocks *before* they update their beliefs about ϕ_t. These perceived values are a function of agents' forecast error on interest rates, which in turn depends on the actual policy rate set by the central bank:

$$\Delta_t = i_t - \hat{\phi}_{t-1} y_t - \rho_v \hat{v}_{t-1} = \left(\phi\left(\beta_t\right) - \hat{\phi}_{t-1} \right) y_t + v_t + \varepsilon_t - \rho_v \hat{v}_{t-1} \tag{13}$$

The following updating equations describe how much of the forecast error gets ascribed to the persistent shock v_t relative to the transitory shock ε_t:

$$\hat{\hat{v}}_t = \rho_v \hat{v}_{t-1} + k_t \Delta_t, \;\; \hat{\hat{\varepsilon}}_t = (1 - k_t)\,\Delta_t, \;\; k_t = \frac{P_{11,t-1}}{P_{11,t-1} + 1}. \tag{14}$$

Here, the coefficient k_t acts like a Kalman gain in the partial updating problem of agents beliefs about the shocks, holding the parameter estimates $\hat{\phi}_{t-1}$ constant.

To find the equilibrium, we have to impose that the observed interest rate equals the interest rate set by the central bank according to equation (2). Substituting (14) into (12), and substituting the resulting expression for y_t into (13) yields:

$$\Delta_t = \frac{\left(\phi_t - \hat{\phi}_t\right)(A_t y_{t-1} + B_t u_t + c_{vt}\rho_v \hat{v}_{t-1}) + v_t - \rho_v \hat{v}_{t-1} + \varepsilon_t}{1 + \left(\hat{\phi}_t - \phi_t\right)(k_t c_{vt} + (1 - k_t)\, c_{\varepsilon t})}. \tag{15}$$

The solution for y_t follows from equations (12) and (14).

References

Clarida, Richard, Jordi Gali, and Mark Gertler (1999). “The science of monetary policy: a new Keynesian perspective,” *Journal of Economic Literature*, vol. 37(4), pp. 1661–1707.

Kreps, David M. (1998). “Anticipated utility and dynamic choice,” *Econometric Society Monographs*, vol. 29, pp. 242–274.