Forecasts and Inflation Reports: An Evaluation

Jon Faust & Eric M. Leeper

June 2005
(revised December 12, 2005)

Abstract

This paper argues that the most important aspect of central bank transparency is communication about the future path of the policy rate. Despite great progress on transparency, communication on this topic remains shrouded in an obscure code. Central banks can enhance transparency by providing, in decreasing order of importance: i) an unconditional forecast of the policy interest rate; ii) an unconditional forecast of goal variables; iii) conditional forecasts of rates and goal variables.

Keywords: DSGE, identification, forecasting
JEL classification:

Federal Reserve Board; faustj@frb.gov (http://e105.org/faustj) and Department of Economics, Indiana University and NBER; eleeper@indiana.edu (http://mypage.iu.edu/~eleeper/).
The paper was prepared for the Sveriges Riksbank conference "Inflation Targeting: Implementation, Communication and Effectiveness," June 11-12, 2005. We have greatly benefited from the comments of Per Jansson, Anders Vredin, Jonathan Wright, participants at the Riksbank conference and participants at Federal Reserve Board and New York Fed seminars. We especially thank Chris Sims for comments including finding a significant error.

1 Introduction

The most distinctive feature of inflation targeting central banks is not their commitment to low and stable inflation: this is now gospel almost everywhere. After more than a decade of experience with inflation targeting (IT), it seems to us that the clearest distinguishing feature of targeting banks is their commitment to the continuing process of improving transparency. While all central banks now issue voluminous public communications, the communication of inflation targeters has continually and rapidly evolved in the direction of clearer and more thorough discussion of policy issues. This process has surely been aided by the fact that inflation targeters have actively invited critical external review and opened their processes to facilitate that review.[1] Our paper is part of this new tradition.

Despite great progress on transparency, we argue that even at IT banks the most important topic of central bank communication remains shrouded in an obscure code. The first question people often ask—until they learn it is futile—of public servants charged with setting the policy interest rate is, "Where are rates going?" This is probably the most useful information central bankers could share with citizens attempting to plan their financial futures. The bulk of the paper demonstrates that current central bank communication about this key question amounts to an obscure code and the paper shows some natural ways to improve transparency.

Here is our basic argument. Communicating the future path of policy is inherently an issue of forecasting. Current inflation reports generally contain one or more *conditional forecasts*—forecasts of output and inflation conditioned on an exogenously specified path of the policy interest rate. We think of these forecasts as a coded message to the public; we argue it is a complex code that is not well-suited to maximizing transparency about where policy and the economy are headed.

While we support these claims more formally below, the key points can be seen in a simple parable. The policymaker is a bus driver and the passengers comprise the public. Full transparency dictates that the bus driver share her private information

[1] Commissioned reports include Svensson (2001), Pagan (2003) and Leeper (2003).

about when the bus will arrive at the destination.[2] The passengers have a basic understanding of the issues—distance, speed limits, and so forth—and form a rough idea of the travel time. The bus driver, who is a regular on this route, has some private information stemming from her specialized knowledge. A passenger inquires about the expected travel time.

Rather than giving an unconditional forecast (the trip usually takes 45 hours), the bus driver reports two conditional forecasts for the principle goal variables: if I drive 60 m.p.h., it will take on average 48 hours, with a negligible probability of a collision; if I drive 80 m.p.h., it will take on average 40 hours, which incorporates delay due to a one percent chance of collision.

How could a passenger decode this message to improve his estimate of typical travel time? If he knows the bus driver's preferences over time versus sustaining a collision, and knows how the probability of collision varies with speed, and knows what delay is typically caused by collisions (which presumably varies with speed at time of collision), then it is straightforward to infer the driver's unconditional forecast of travel time. Otherwise, it may be unclear how the conditional forecasts are of *any* value to the passenger hoping to improve his own rough estimate of arrival time.

To complete the bus story, we might imagine the persistent passenger asking for an unconditional forecast. The bus driver responds that if she gave an unconditional forecast, the passengers might try to hold her to it; this could make it difficult for her to react appropriately to unexpected changes in driving conditions. True enough.

For most of the paper we simply discuss how to improve transparency, leaving aside whether this is desirable. In the final section, we return to the question of whether clearer communication would be a good thing.

[2] Andy Levin notes that we are taking it as given that the public knows where the bus is going, which some might equate with stating an inflation target.

2 Transparency and forecasts

We assert three transparency principles. First, absent competing interests, central banks should strive to shed as much of their private information as is reasonably possible. The least ambiguous foundation for this principle is that, as Chairman Greenspan (2002, pp. 5–6) of the Fed argued, "Openness is an obligation of a central bank in a free and democratic society. . . ."[3] Transparency might also improve policy outcomes, but as the transparency literature makes clear, in a distorted economy there are many reasons why it might be welfare enhancing for the central bank deliberately to surprise the public.[4]

Second, a main focus of central bank transparency efforts should be to shed private information about the future course of policy. The central bank's main social policy mandate is to conduct monetary policy, and this gives the central bank a particular responsibility to share information about the course of policy.[5]

Third, in pursuing transparency a central bank using an interest rate instrument should be transparent about the future path of instrument. This principle follows more or less directly from the first two.

Some might argue that it is sufficient for the central bank to be transparent about the goal variables of policy. Perhaps, for example, communication could focus only on inflation and the output gap without great loss.

The analogous argument in the bus parable is probably correct. It is important that the bus driver communicate about goals such as travel time, smoothness of the ride, and safety. It is not nearly as important that she communicate about the instruments of policy—how the gas pedal and steering wheel will be manipulated to attain the goals.

In this respect, the bus example misses an important element of the monetary

[3] See also the remarks by Issing (2004).

[4] Faust and Henderson (2004) develop this point more fully.

[5] Further, the primary *inherent* informational advantage possessed by the central bank is about the future course of policy. Of course, the central bank may have private information about the economy more generally—even as distinct from that which flows from a better knowledge of the course of policy. This information should also be shared.

policy case. Individual welfare is strongly affected by the path of the interest rate instrument through direct channels not captured by the conventional goal variables. For example, changes in nominal interest rates can lead to important wealth redistributions. These may largely net out in aggregate welfare calculations, and economists often do not emphasize them in analyzing aggregate policy. Such issues should not be neglected, however, when considering what a central bank has a responsibility to communicate to the public it serves. A person in a very secure job but facing an interest-sensitive decision, such as securing a fixed-rate mortgage, may care much more about the short-run path of interest rates than about the short-run path of the inflation and output gaps.

2.1 Why forecasts are central to transparency

The transparency principles present the central bank with a difficult communication problem. The central bank's information set is large and amorphous, and any communication strategy must distill this information. Here we begin the formal analysis of information and illustrate why certain forecasts are the natural way to encode central bank information for purposes of public communication. While we have not seen these points collected in this way, they are present in many places.[6]

Under any formulation, the policy process maps current information into a policy choice:

$$i_t = \phi(Y_{t-1}), \tag{1}$$

where Y_t is the central bank information set at t, and at this point the time subscript may measure days or minutes. In the realistic case, Y is high dimensional and cannot be fully characterized; ϕ is imperfectly known, even to the policymakers. At the policy meeting, the policy board has a wide ranging discussion—explores Y_{t-1}—and ultimately chooses a policy—evaluates ϕ.

The goal in designing a communication policy is to reduce or condense the char-

[6] For example, Sims 1980; Bernanke, et al. 1999; Leeper and Zha 2003; Svensson 2004, Woodford 2004.

acterization of the current information and the policy mapping to something that can be communicated effectively. Two conventional assumptions simplify the problem considerably.

First, assume that policy can approximately be seen as period-by-period minimization of some expected loss function:

$$i_t = \arg\max_i E\left[L(Y_t(i))|Y_{t-1}\right] = \int L(Y)\xi(Y_t(i)|Y_{t-1})dY. \tag{2}$$

The second expression expands the definition of the conditional expectation in somewhat short-hand notation, assuming that $\xi(Y(i)|Y_{t-1})$ is the joint conditional density of Y_t when the history is Y_{t-1} and the policy rate is $i_t = i$.[7]

Since the density $\xi(Y_t(i)|Y_{t-1})$ may be thought of as a forecast equation, (2) emphasizes the tautology, through the meaning of *expected*, that expected loss minimization can be written in terms of forecasts.

The second assumption is that the loss function itself can be approximated as a function of a much condensed information set:

$$L(Y_t) \approx l(Z_t), \tag{3}$$

where $Z_t = \Xi(Y_t)$ and is much lower in dimension than Y_t. In the familiar case, Ξ picks out the two gaps. Now we can write,[8]

$$i_t \approx \text{argmax}_i = \int l(Z_t)\xi(Z_t(i)|Y_{t-1})dZ. \tag{4}$$

We have now greatly simplified the communication problem: to be transparent about the future course of policy, the central bank only needs to communicate its loss function over Z and its forecasts for this low-dimensional variable. The only additional note at this level of generality is that Z should include items of great importance to individual welfare, which includes the policy interest rate.

[7] Note that we are not necessarily assuming that L is a true social welfare function, so this expression covers a wide range of cases. In particular, this covers the formulation that Svensson has emphasized in a number of papers (e.g, 2004). It also summarizes a conventional approach to making policy in intractable problems as described in Faust (2005), in which case, L is an *ad hoc* judgmentally derived *scoring function*.

[8] Note that the forecast is still conditioned on the full information set.

2.2 Current Inflation Reports

As is conventional in this literature, we call forecasts conditional on central bank information at the time of the forecast and on a future path for the policy rate *conditional*. *Unconditional forecasts* are conditional only on information available to the central bank at the time of the forecast and not on any counterfactual paths for variables in the forecast period.[9]

Reporting of forecasts is a topic of rapid development in inflation reports. Due to the rapid developments, we can at best present a representative snapshot of the range of interest rate assumptions and the diversity in how the assumptions are treated. Initially, the main forecast was conditional and typically premised on a constant policy rate. There was no pretense that policymakers regarded the flat path as plausible. The European Central Bank continues in this tradition [European Central Bank (2005)].

More recently, emphasis has shifted to a forecast conditioned on policy following a *market path* implied by the term structure of forward interest rates. The Bank of England (BOE) has always produced a constant rate and a market rate projection. Initially, the constant rate was emphasized, but the central projection has been based on the market path since the August 2004 report. Sveriges Riksbank (Sweden) converted to market rates in the second half of 2005, and a constant rate projection is still reported.

Two central banks appear to condition on interest rate paths that are chosen to achieve their policy objectives. Since the beginning of 2005, the Norges Bank (Norway) computes projections conditional on an interest rate path explicitly endorsed by the Executive Board as striking a "reasonable balance between the objectives of monetary policy" [Norges Bank (2005), p. 5]. For many years, the Reserve Bank

[9] Some confusion could arise because there are other variables than the policy rate that are taken as given in certain steps of the forecast process at central banks. For example, often the paths of certain hard-to-forecast variables are set judgmentally and taken as given in the rest of the forecast procedure. So long as the judgmental paths are some sort of "best guess" and not deliberately counterfactual, the fact that they are taken as given during some forecast steps does not render them *conditional* in the sense we are using it or *exogenous* in any standard sense of the term. This claim about conditionality relies on the fact that we are only considering first moments.

of New Zealand (RBNZ) has conditioned its projections on an interest rate path its Governor deems is consistent with achieving a forecasted path of inflation within the target bands over the medium term [for example, Reserve Bank of New Zealand (2005)].

Although there has been a good deal of argument that moving to the market-based assumption is a significant improvement, we will argue that, at least from the standpoint of economically sophisticated agents, the information content of *any* conditional forecast is essentially the same. To the extent that the interest rate paths used by Norway and New Zealand reflect the policymakers' own views of where policy is headed, however, their approaches are close to the unconditional forecasts we have in mind.

2.3 Consumers of forecast information

Because this paper mainly concerns communication to the economically sophisticated public (ESP), it is useful to distinguish three consumers of forecast information.

The ESP efficiently extracts all relevant information from central bank communication. In contrast, the typical person is someone who, absent further digestion or pedagogy, cannot fully appreciate the implications of all available central bank information.[10] We leave to the conclusion some discussion of communication to this group.

This paper is also not about communication between the central bank staff and policy board. Under the assumptions that the Board is sophisticated and that it has access to the full information set upon which the forecast is based, the forecast—no matter what its properties—has no clear implications for the policymaking step. So long as the policymakers are aware of the model's shortcomings, bad properties of the model or forecast need not show through to policy.[11]

[10] Markets and the media aggregate and reveal information; thus, individuals may acquire information that they could not individually have extracted from raw central bank communication.

[11] Here our approach differs from, for example, Woodford (2005) who argues that flaws in the

3 The information content of forecasts

We begin our argument that conventional conditional forecasts are an obscure way to communicate the policy outlook with a simple example. The economy is described by a Phillips curve and inflation equation. The essential feature for our purpose is that the central bank bases policy on some private information. After receiving a private signal about the state of the economy, the central bank sets the policy rate to minimize a quadratic loss function in an output gap and inflation.[12]

In this section, we assume that the central bank's information set is a superset of that of the public. Thus, the public uses any forecast announcement to infer the central bank's information and update its own view of where policy and the economy are going. For simplicity, we focus only on the point forecast, leaving the issue of higher moments and fan charts aside.

3.1 The model

The Phillips curve and inflation equation are:

$$x_t = -ai_t + b(\pi_t - \pi_t^e) + u_{xt} + w_{xt} \quad (5)$$

$$\pi_t = -ci_t + u_{\pi t} + w_{\pi t} \quad (6)$$

where x_t and π_t are publicly observed measures of the output gap and inflation at t, respectively, and π_t^e is the public's expectation of inflation formed at the end of $t-1$ and is observed by the central bank. Everyone knows these equations and their coefficients.

The u and w variables are mean zero, finite variance, and serially uncorrelated shocks. These are not observed directly by anyone. The two w shocks are standard

forecast process could have implications for policy. This case arises if the Board, loosely speaking, *takes the model literally*. We instead assume that the Board has access to the full information set including the model, and *takes the model for what its worth*, discounting aspects that are not helpful in formulating optimal policy. Something in between these two assumptions is probably correct.

[12] In simple models, there is always a solution to such communication problems: tell people your information. To be clear, we are not studying cases in which the central bank has any tendency to lie to the public.

economic shocks. The u_π and u_x represent a gap between the measured inflation and output gaps and those that are relevant for welfare. We think of these as measurement error in a broad sense—that is, as the gap between announced data and ideal measures. The separate role of the u's and w's can be seen in the social loss function:

$$E_t \frac{1}{2}\left[(\pi_t - u_{\pi t})^2 + \lambda(x_t - u_{xt})^2\right]. \tag{7}$$

The bliss points for the true inflation and output gap measures are normalized to be zero. Welfare is a function of the w's, but not the u's. Before setting policy each period, the bank staff's efforts give rise to an unbiased signal of each of the four shocks: $\hat{v} = \{\hat{u}_{xt}, \hat{u}_{\pi t}, \hat{w}_{xt}, \hat{w}_{\pi t}\}$.

The timing of events is as follows. At the end of $t-1$ the public forms its expectation, π_t^e. At the beginning of t, the central bank receives the signal, $\hat{v}_t$, and the central bank sets i_t optimally, taking π_t^e as given. Then the shocks, v_t, are realized and x_t and π_t are determined. Next, the cycle begins again with the public forming an expectation for $t+1$.

This model has a simple linear-quadratic structure, and the solution is conventional (see the Appendix for details). The relevant features for our purposes are as follows. In all periods, $\pi_t^e = 0$: at the end of $t-1$ neither the public nor the central bank have any information to suggest that non-zero inflation will be optimal. The central bank reaction function will be a linear function of central bank news, and, in particular,

$$i_t = \alpha_1 \hat{u}_{\pi t} + \alpha_2 \hat{w}_{\pi t} + \alpha_3 \hat{w}_{xt}. \tag{8}$$

with, $\alpha_1 \neq \alpha_2 \neq \alpha_3$.[13]

[13] The basic idea of the solution is that the central bank wants to optimally smooth the shocks that affect welfare. Obviously, it reacts to both of the w shocks. While u_x has no effect on welfare and is ignored, u_π indirectly affects welfare through its effect on the gap via the inflation surprise in the Phillips curve.

3.2 Transparency through forecasts

After the signal arrives and before the policy rate is observed, the central bank publishes an inflation report containing one or more forecasts of the economy. We consider what aspects of the central bank's private signal can be encoded in various sorts of forecasts.[14]

Consider a forecast conditional on the signal, $\hat{v}$, and on a policy assumption, $i_t = i^c$. The forecast will be given by the conditional expectation of (5-6):

$$x_t^c = -(a+bc)i^c + bc(\hat{u}_{\pi t} + \hat{w}_{\pi t}) + \hat{u}_{xt} + \hat{w}_{xt} \quad (9)$$

$$\pi_t^c = -ci^c + \hat{u}_{\pi t} + \hat{w}_{\pi t}. \quad (10)$$

In the x equation we have substituted for π to write the system in terms of policy and shocks.

If the coefficients and policy assumption are known, it is clear from (9) and (10) that announcing a conditional forecast is equivalent to simply announcing the sum of the two errors in the structural equations—$(\hat{u}_x + \hat{w}_x)$ and $(\hat{u}_\pi + \hat{w}_\pi)$. While this result is specific to the model, the general point is that the conditional forecast reveals some, but not all of the central bank's information.

To make this point more concrete, consider how the public can use the conditional forecast to infer the central bank's unconditional forecast. Denote the central bank unconditional forecast with a superscript *cb* and the public's updated forecast after seeing the conditional forecast with a *p*:

$$i_t^{cb} = \alpha_1 \hat{u}_{\pi t} + \alpha_2 \hat{w}_{\pi t} + \alpha_3 \hat{w}_{xt} \quad (11)$$

$$i_t^{p} = \alpha_1 E[u_{\pi t} | \hat{u_\pi} + \hat{w}_\pi] + \alpha_2 E[w_{\pi t} | \hat{u_\pi} + \hat{w}_\pi] + \alpha_3 E[w_{xt} | \hat{u_x} + \hat{w}_x] \quad (12)$$

Because the public knows only the sum of the two shocks in each equation, it cannot precisely extract the central bank's expectation from the conditional forecast.

[14] One might wonder why the public would bother to read the inflation report: the information cannot affect outcomes in this model. We have deliberately imposed this simplicity. Still, it is fully consistent with the model for us to imagine that, after reading the inflation report, agents make decisions that affect individual, but not social, welfare.

The difference in policy expectations will lead the public and central bank forecasts of the other variables to diverge in a very simple manner. To see this, difference the expressions (9) and (10) for forecasts based on two different conditioning paths, i^c and $\tilde{i}^c$, giving

$$\tilde{x}_t^c - x^c = -(a+bc)(\tilde{i}^c - i^c) \tag{13}$$

$$\tilde{\pi}_t^c - \pi^c = -c(\tilde{i}^c - i^c) \tag{14}$$

Since these equations are of the same form for any forecasted variable, it is useful to define z as the typical forecasted variable, and write the generic equation as,

$$\tilde{z}^c - z^c = \gamma_z(\tilde{i}^c - i^c) \tag{15}$$

where γ_z is the coefficient on i in the reduced-form equation for z (treating i and the shocks as exogenous). This equation allows us to draw several strong conclusions, which carryover to the richer dynamic and nonlinear case considered below.

First, any single conditional forecast contains the same information. This is clear from the fact that the difference of any two such forecasts in (15) involves only γ and the policy assumptions, both of which are assumed known. This also implies that, so long as agents know γ, conditional forecasts after the first contain no marginal information.

Second, if agents are uninformed about γ, a second forecast would reveal γ; additional conditional forecasts would reveal no marginal information.

Third, after the public updates its forecast upon seeing the inflation report, the public and central bank forecasts for z will differ to the extent that their policy rate forecasts differ. This can be seen in the version of (15) based on the public and central bank forecast of rates,

$$z^p - z^{cb} = \gamma_z(i^p - i^{cb}) \tag{16}$$

We now show that these results follow generally.

3.3 The dynamic case

Suppose that each period the central bank forms an unconditional forecast denoted, as above, by *cb*. We place no restrictions on how this is formed. Conditional forecasts are then formed consistent with Leeper and Zha's (2003) formulation of modest policy interventions—that is, they are treated as a sequence of hypothesized policy shocks of a modest magnitude. We assume that in a modest neighborhood of the unconditional forecast, the responses of all relevant variables, including the policy rate, to the policy shock are linear in the shock and well approximated by the conventionally defined impulse responses of the variables to a policy shock. These responses may be different for different unconditional forecasts. Thus, we are principally assuming that the model is sufficiently smooth (at least in the policy variables) to make local linear approximation useful.

Under these assumptions, a few lines of algebra get us to the dynamic analog of the static results. Suppose we condition on a policy rate path of $(i^c_t, \ldots, i^c_{t+h})$. The conditional forecast for i and typical variable z can be written as a deviation from the unconditional forecast using the assumed impulse responses:

$$z^c_{t+h} = z^{cb}_{t+h} + \sum_{j=0}^{h} d_i \varepsilon^c_{t+h-j} \tag{17}$$

$$i^c_{t+h} = i^{cb}_{t+h} + \sum_{j=0}^{h} g_i \varepsilon^c_{t+h-j} \tag{18}$$

where the impulse responses of z and i to policy shocks are given by the d_j's and g_j's, respectively. The second equation implicitly defines the policy shocks, ε^c_s, that are consistent with the conditioning path for rates. In matrix notation, these equations can be written,

$$z^c - z^{cb} = D\varepsilon^c \tag{19}$$

$$i^c - i^{cb} = G\varepsilon^c \tag{20}$$

where $z^c = (z^c_{t+0}, \ldots, z^c_{t+H})'$ and the other superscripted variables are analogously defined. Finally, $D = Q(b)$, where $Q(b)$ is the lower triangular Toeplitz matrix with

b_i on the i^{th} subdiagonal; $G = Q(g)$ is analogously defined. Thus,

$$z^c - z^{cb} = \Gamma(i^c - i^{cb}) \tag{21}$$

where $\Gamma = BG^{-1}$. As above, we can difference this expression for two conditional forecasts:

$$z^c - \tilde{z}^c = \Gamma(i^c - \tilde{i}^c) \tag{22}$$

Of course, (21) and (22) mirror (16) and (15). Thus, the three conclusions listed above follow directly.[15] We have lost something significant, however, in going to the more general case: we no longer have an explicit expression for how the public should update its policy expectation to form i^c. The three conclusions follow without this, but some further comment is warranted.

In the simple case, an explicit expression, (12), for how the public should form i^p could be derived because the public knew the central bank reaction function and had a complete characterization of how the news in the conditional forecast should affect the arguments of that function. Even without additional structure, it is possible to characterize some results about the information content of conditional forecasts, however. We know that the public's updated forecast must satisfy (21) for each variable z. Stack (21) for, say, N different z variables announced by the central bank to get,

$$\begin{bmatrix} z_1^p \\ \vdots \\ z_N^p \end{bmatrix} = \begin{bmatrix} z_1^c \\ \vdots \\ z_N^c \end{bmatrix} + \begin{bmatrix} \Gamma_1 \\ \vdots \\ \Gamma_N \end{bmatrix} \begin{bmatrix} i^p - i^c \end{bmatrix} \tag{23}$$

Treating all the superscript p variables (z's and i's) as unknown, this is a set of Nh equations in $(N+1)h$ unknowns. These form a set of restrictions that any mutually consistent choice of $(i^p, z_1^p, \ldots, z_N^p)$ must satisfy. Because of the triangular form of Γ, a more specific way of characterizing the restrictions is that beginning with the first forecast horizon, these equations leave one degree of freedom in choosing the values of $N+1$ variables at each horizon. If the central bank announces two variables,

[15] We will show that (15) can be solved for Γ below.

π and x, then there is one degree of freedom in choosing (π, x, i) triplets at each horizon. As in reality, we do not have an explicit reaction function or objective function for the central bank. The public could, however, use its general knowledge of the central bank's goals to refine its choice among admissible (π, x, i) triplets.

This essentially completes the theory part of our support for the claim that sharing a conditional forecast is an obscure way of encoding the central bank's view of the likely path of policy and the economy. The conditional forecast generates the set of restrictions, (23), that should be satisfied by the public's updated forecast. The remaining step of choosing a forecast consistent with these restrictions will depend on the particulars of the public's knowledge of the central bank's objective function.

We can take the argument a bit further by showing that the coefficients, Γ, which play a key role in the discussion, are themselves a commingling of more standard parameters. Using the Toeplitz form of Γ, (21) can be written as,

$$z^c_{t+h} - z^{cb}_{t+h} = \sum_{j=0}^{h} \gamma_{zj}(i^c_{t+j} - i^{cb}_{t+j}) \tag{24}$$

with $\gamma_z = (\gamma_{z0}, \gamma_{z1}, \ldots, \gamma_{zh})'$ and $\Gamma_z = Q(\gamma_z)$. Equations (24) and (18) are identical in form, but (24) is written in terms of the deviations between the unconditional and conditional paths for policy, whereas (18) is written in terms of the implied policy shocks.

The similarity of these two forms sometimes leads to confusion, with the γ's in (24) being interpreted as the impulse response of z to a policy shock. Of course, the d's form this response. A bit of algebra reveals that the γ's are commingling of the impulse responses of z and i to the policy shock. We will call the collected γ's the Γ-response to distinguish is from the conventional impulse response. As noted above, two conditional forecasts allows the public to infer γ, but this is not a parameter that is easily interpreted.

3.4 Discussion

The previous section was quite general in that it put essentially no restrictions on how the unconditional forecast was generated. The results flowed from assuming that for modest policy interventions, the response to a policy shock was approximately given by a set of linear impulse responses for the variables in question.

The conditions of this analysis need not be met in practice. In our view, however, if they are not met our main point is likely to be even stronger. Suppose that the responses of the economy to different modest interventions differ in importantly nonlinear ways. In this case it is difficult to see how one or two forecasts based on arbitrary paths are of much value in informing the public about what would happen under the particular path given by the unconditional.

Recently, a number of economists have argued in favor of announcing conditional forecasts conditioned on some publicly observed measure of the public's expectation of policy such as a path implied by futures market prices. As noted above, some central banks have moved to emphasizing such forecasts. Our analysis sheds some light on the merits of this practice.

So long as the central bank is choosing among different conditioning assumptions, *each of which would constitute modest policy interventions* relative to the unconditional, then the market-based path has no advantage over any other. Conditioning on a publicly available measure of the market's own forecast does not reveal additional information of the central bank.

One advantage of the market-based path is that it is likely to be in the neighborhood of the central bank's own unconditional path for policy, and is, thus, likely to constitute a modest intervention. At times, other conventional assumptions such as a constant interest rate may be wildly counterfactual and the approximate linearity of the response of the economy may be implausible. Indeed, a symptom of this implausibility may be the breakdown of conventional means for generating a model-based forecasts. In such cases, the conditional forecast will be difficult to form and interpret. Thus, while the market-based path is not inherently more infor-

mative than forecasts based on any other modest-intervention path, it may represent a convenient way to pick a modest intervention path.

4 Forecasts and transparency in practice

In the previous section, we assumed that the central bank's information is a superset of the public's and that both the central bank (CB) and the sophisticated public use all information efficiently. Suppose that there is a possibility that the CB does not use information efficiently—that is, there is some chance that the forecast is systematically deficient in some respect.

The possibility that CBs are imperfect forecasters adds a potentially overwhelming element to the argument that conditional forecasts are an obscure way to encode important information. Define z^0 as the public's forecast of z before seeing the inflation report. Under the efficiency assumptions of the previous section, we can write,

$$\begin{aligned} z^0 - z^c &= \Phi_1 + \Phi_2 \\ \Phi_1 &= z^0 - z^{cb} \\ \Phi_2 &= \Gamma(i^c - i^{cb}) \end{aligned}$$

The difference between the public's unconditional forecast before seeing the report and the conditional forecast in the report is due to two terms: Φ_2 is due to the forecast being conditional; Φ_1 is the difference between the public's unconditional and the CBs. Since Φ_1 is purely due to the CBs superior information, the public simply wants to remove its best guess of the second term in order to deduce z^{cb}.

Suppose that CB's forecast is not fully efficient and that the CB makes systematic errors—relative to its information set. Under this assumption, some portion of Φ_1 is systematic error in the CBs forecast, which can be predicted using the information set of the public:

$$\Phi_1' = (\text{CB inefficiency}) + (\text{CB superior info})$$

In deciding how to update its forecast based on the conditional forecast of the CB, the public now has to remove the effects of conditioning, but then must decide what portion of the remainder is due to information superiority of the CB and which portion is due to inefficiency of the CB.[16]

It is important to emphasize that for modest policy interventions, Φ_2 may be fairly small in magnitude: $(i^c - i^{cb})$ may be small and the effect of such small deviations on the economy, especially at short horizons, is also likely to be small. Thus, the term in Φ_1' due to inefficiency could easily dwarf Φ_2 in magnitude.

4.1 Evaluating efficiency based on conditional forecasts

4.1.1 Background

How should the public go about deciding if the CBs unconditional forecast is efficient? If the CB only publishes conditional forecasts, this is a difficult problem. Until recently, the best academic practice in analyzing CB forecasts seems to have been to simply analyze the conditional forecasts as if they were unconditional, noting that this is an unpleasant compromise. By this standard, some CB forecasts have been found to be pretty good. Romer and Romer's (2000) work shows that the Fed's Greenbook forecast outperforms conventional benchmarks and Sims (2002) reaches a similar conclusion.[17]

Are these positive results? It is not clear that a forecast based on deliberately counterfactual policy assumptions should perform well as an unconditional forecast.

Based on assumptions and results like those in the previous section, Faust and Wright (2005) derive one way to test efficiency of an underlying unconditional forecast based only on data from conditional forecasts. The test starts with a row of (21):

$$z^c_{t+h} - z^{cb}_{t+h} = \sum_{j=0}^{h} \gamma_{zh}(i^c_{t+h} - i^{cb}_{t+h}).$$

[16] We note that if the CB uses information inefficiently in forecasting the policy rate, we could similarly decompose Φ_2. We ignore this point for simplicity.

[17] For a slightly different result, see Faust, et al. 2004.

If the CB forecast is efficient,

$$\begin{aligned} z_{t+h} &= z^{cb}_{t+h} + \nu_{z,t+h} \\ i_{t+h} &= i^{cb}_{t+h} + \nu_{i,t+h} \end{aligned}$$

where the ν's are unforecastable by information known at t. Substitute out the CB forecast using these efficiency conditions:

$$z^{c}_{t+h} - z_{t+h} = (\sum_{j=0}^{h} \gamma_{zh}(i^{c}_{t+h} - i_{t+h})) + \nu_{t+h} \tag{25}$$

where ν_{t+h} gathers up the various ν terms. The error is correlated with $(i^{c}_{t+h} - i_{t+h})$, but not with any information available at t. Thus, we can estimate the equation by instrumental variables using instruments in the time t information set. We can test the efficiency of the CB forecast by augmenting the equation with predictors, w_t, known at t. If the *unconditional* CB forecast is efficient, no such w's should explain the error in this instrumented *conditional* forecast error model.

A byproduct of this procedure is a consistent estimate of the γ's. Thus, in the theory section we showed that two conditional forecasts allows one to calculate the γ's. Now we add that, so long as the γ's are constant through time, a time series of conditional forecasts at different points in time allows one to estimate the γ's.

The econometric approach and Faust and Wright (2005) allows us to advance the critical study of conditional forecasts. From the standpoint of this paper, however, the most important lessons from Faust and Wright regards CB transparency.

First, these approaches are very complex: properly implementing these procedures gets into complicated issues of instrumental variables estimation and the analysis of weak instrument econometrics. Forcing the public to use such techniques to decode the inflation reports cannot advance the cause of transparency.

Second, due to short samples, these approaches will be of limited value to the public. A back-of-the-envelope calculation may serve to illustrate this point. Suppose that 33 percent of the variance of the forecast error of the CB is predictable using publicly available information. This strikes us as an implausibly large value

but will illustrate the point. A typical test statistic for forecast efficiency involves a regression of the forecast error on variables known at the time of the forecast. The statistic is approximately sample size times the marginal R^2 from adding variables known at the time of the forecast. Under the null hypothesis of forecast efficiency, the statistic will be distributed as χ^2 with degrees of freedom equal to the number of predictors added. With 4 years of quarterly inflation reports (12 observations) and a marginal R^2 of 1/3, the p-value of this test will be 0.27 if based on three predictors and 0.14 and 0.05 with 2 and 1 predictors, respectively. Even if the CB is a very inefficient forecaster, evidence of this fact accumulates rather slowly. If the inefficiency arises due to the combined effect of many sources of information, the degrees of freedom will be large relative to the entire current history of inflation reports and inferences will be imprecise.

4.1.2 Detail on methods

Since the BOE presents two conditional forecasts, our simplest results come from computing the implied γ's. The remainder of the results come from instrumental variables estimation of equations of the form (25). Given this framework, three choices remain. First, we can estimate these equations separately for each horizon or jointly as a system. Using the single equation approach we do not impose that the γ's should be the same in each equation. In a system, we can impose this restriction. We report both types of estimates.

Second, we need to choose instruments for $i^c - i$, the *ex post* difference between the conditioning path and the outcome for the policy rate. The instruments need to be in the information set at the time the forecast is made. Because the conditioning path is deliberately counterfactual, a natural choice of instrument is the difference between the conditioning path and a good forecast of the policy rate at the time of the forecast. Of course, the forward interest rate implied by the term structure of interest rates provides a one good forecast of rates, so it is a natural instrument.

Third, we must choose alternate predictors, w. In most forecasting work, the

forecast itself is used as a predictor of the error. This is our main source of alternative predictors. Of course, because of the simultaneity between interest rates and inflation and output, anything that predicts the forecast error in output or inflation might also help predict the *ex post* difference between the interest rate and the conditioning path. Thus, we should also include our w's as potential instruments.

Our efficiency test, then, amounts to a test of whether the instruments explain the forecast error *in any other way* other than through their ability to prediction $i^c - i$. So long as there are more instruments than free parameters, γ, this can be tested. In the linear IV case, this is tested using the standard Sargan test. As discussed above, the test statistic is essentially the sample size marginal R^2 from allowing the instruments to enter directly rather than only entering as instruments. In the GMM context with cross-equation restrictions, the generalization of the Sargan test is Hansen's J-test.

We report both estimates of γ and the associated J-tests, which we interpret as a test of forecast efficiency. We report GMM estimates of the system from the continuous updating criterion function and conventional asymptotically justified standard errors based on the derivatives of the criterion function.

4.2 Data

We use the GDP and inflation forecast data published by the Riksbank and BOE. The Greenbook data were obtained internally at the Federal Reserve Board, but equivalent data are published (with the 5-year delay) on the Philadelphia Fed website.[18] The vintage GDP data for Sweden were provided by the Riksbank. Some of the vintage data for the U.K. are on the BOE website; we augmented this with data of our own construction. For the Riksbank and BOE inflation measures, revisions are not an issue. The U.S. measure we study is the GDP deflator and is revised. We constructed the required vintage data for the U.S. ourselves.[19]

[18] Actually, some of the data such as the interest rate assumption data are not yet up, but will soon be.

[19] One could use the data on the Philadelphia Fed web site, but the vintages are not timed to correspond with Greenbook forecasts. Our data are precisely timed to correspond to information

For the BOE we study the quarterly forecasts from 1998:1 through the current. For the Riksbank, we study 2001:1 through the current. For the Greenbook, we study forecasts from Oct. 1988 through the last Greenbook in 1999.

For the BOE we have two conditional forecasts. One is based on a constant policy rate assumption (denoted bar); the other on a policy rate path implied by futures markets (denoted mkt). For the Riksbank, we have only a constant rate (bar) forecast. For the Greenbook over this sample period, the assumption varies and was sometimes constant and other times not. It was never simply a rate implied by forward markets.

Some of our work relies on forward estimates of the policy interest rate implied by financial markets. The BOE and Riksbank publish this information in the inflation reports. For the BOE it is the rate used in the conditioning path of the market-based forecast. We constructed the variable from Federal Funds rate and eurodollar futures data for the U.S.[20]

Two additional issues must be mentioned. The BOE and Riksbank publish forecasts for annual percent changes. We have converted these to quarterly logarithmic changes computed as $400\log(x_t/x_{t-1})$. Some of our results are reported for annual changes, calculated as $100\log(x_t/x_{t-4})$.

Finally, with all forecast work involving vintage data, there is a question about what to treat as the final data in computing forecast errors. People differ on the appropriate approach.[21] In this paper, for any forecast done in quarter t we use the data as they stood about 2 years later in computing the forecast errors. Thus, when our work involves computing forecast errors, this shortens the available sample of forecasts.

used in constructing the Greenbook forecast.

[20] The approach to construction is described in Faust, Rogers, Wright and Wang, (2003).

[21] For a discussion of this issue, see Faust, et al. (2005).

4.3 Results

CBs and others have performed forecast evaluations before. Each of these banks performs regular analysis of the forecasts and the inflation targeters commission external reviews of their work, including the forecasts. Work by the Bank of England (2004) shows that their inflation forecast is pretty good by conventional standards, but the GDP forecast is not as good. Several authors have pointed out that the Riksbank and BOE forecasts have somewhat puzzling features when viewed as conditional forecasts (Faust and Henderson, 2004; Leeper, 2003; Goodhart, 2005). In particular, over most of the history of the published forecasts, the inflation forecast conditioned on a constant policy path returned to the inflation target at the relevant horizon stated by the bank. Given the inflation targeting framework, this would seem to imply that no change in rates would be predicted. In practice, rates were changed regularly, and the changes were quite predictable.

Since our goal in this paper is to critique the use of conditional forecasts as communication devices and not thoroughly to critique the forecasts themselves, we only present some illustrative results to make two points: it would be reasonable for the public to question the efficiency of CB forecasts; in the face of such questions the interpretation of conditional forecasts is muddied considerably.

4.3.1 BOE Γ-responses implied by pairs of conditional forecasts

One might wonder what the Γ-responses ought to look like. As a frame of reference, we take a conventional set of impulse responses to a policy shock for the U.S. as reported in (Faust and Rogers, 2003). The impulse responses and the implied Γ responses are shown in Table 1. The policy shock is normalized to have a unit effect on rates initially which rises for a period and then decays. The response of output is zero on impact by assumption and then monotonically increases in magnitude over the 5 periods shown. While the effect of the shock on output is largest at lag 5 this effect is attributable to all the implied policy shocks from time 0 through 5; thus, the Γ response at lag 5 is smaller than at lag 4.

The Γ-responses implied by pairs of BOE forecasts in each inflation report are shown in Table 2. The main impression one gets is that the implied responses are quite variable across inflation reports. Further, across lags in a given row, the responses often fluctuate in sign and appear explosive. In short, it is difficult to imagine that the public could learn something valuable about the CBs understanding of policy effects from these responses.

As Archer (2004) emphasizes, in a judgmental forecasting framework, it is very difficult to create a set of forecasts that have the sort of structure and internal consistency that would allow us to derive the sort of inferences we are attempting here. This may be a further reflection of the point emphasized by Goodhart (2005) that the forecast seems to imply that no change in policy would be required, but in fact policy changed regularly and in a predictable manner.

4.3.2 Estimated Γ-responses and efficiency tests

Estimated Γ-responses and efficiency tests are presented in Table 3. For the inflation and GDP forecasts of each bank, we present estimates from both single-equation and systems estimates. For the single-equation estimates for horizon h, we use the difference between the forward rate and conditioning path for horizons 0 through h and the forecast at horizon h giving rise to the left-hand side forecast error. For the system of equations 0 through h, we use only the 0 through h interest differentials. Although not strictly necessary, we include a constant in each equation.

Our reasoning about small samples and the Table 2 results for the BOE suggest that erratic γ estimates might be expected. This is what we find. Across different rows of the table, the point estimates differ greatly. On any row, the estimates often oscillate in a manner that may indicate explosive behavior. Generally the standard errors of the point estimates are at least as large as the point estimates.

The J-tests give a bit of evidence against forecast efficiency, mainly for the GDP portions of the forecast. The p-values are sometimes near or below 10 percent, for example. As noted above, even if the forecasts are quite inefficient, we would only

expect weak evidence from the small sample that we currently have on hand.

We include estimates for the Federal Reserve's Greenbook forecast as a basis for comparison. In this case we have 85 observations and might expect better results. In general, the point estimates vary much less across rows and the standard errors are much smaller. It is still true, however, that many of the coefficients are not statistically different from zero at conventional significance levels.

The overall conclusion we take from this is that evidence on the γ's and on forecast efficiency accumulates slowly. Thus, if the public believes that there may be some forecast inefficiency or needs to infer γ to interpret the CB's conditional forecast, it may take decades for inferences to be anything but very hazy.

5 Discussion

There seems to have been an aura that conditional forecasts as published in some inflation reports have played an important role in communicating about policy. Some informal reasoning goes like this: if the forecast under fixed policy shows inflation going up, that signals that the CB is likely to tighten. The Riksbank even embedded this logic in a rule of thumb (Sveriges Riksbank, 2003).[22] As Faust and Henderson (2005) emphasize, even this simple reasoning does not hold with multiple goals as under flexible inflation targeting. In this case, one would need a statement about real activity and some information about how to balance the goals. To correctly reason about the trade-off requires analyses of the impulse responses of both output and inflation to a policy shock, leaving us with precisely the messy sort of analysis we provide above.

Our view is probably clear, but in the interests of transparency we will state it unambiguously. Conditional forecasts as they have generally been published to date in some inflation reports have been of little, if any, value. It would be difficult to directly verify the value of the forecasts, positive or negative. Given the properties cited above, it seems to us that the forecasts have had as much potential to confuse

[22] With the new emphasis on the market-rate forecast, this rule of thumb has been de-emphasized.

as to clarify.[23]

One lesson from this analysis might be that the CBs should improve the quality of the conditional forecasts. The suggestion to make better forecasts is, of course, uncontroversial. From the standpoint of transparency, however, improving the conditional forecasts represents an attempt to improve the code used to encode the message. We take a different lesson. Abandon the code and attempt to communicate directly: state a view about where policy and the economy are headed.

To be a bit more concrete, we suggest that to maximize transparency, the CB should publish an unconditional forecast of the policy rate and goal variables. For simplicity, we have limited the analysis in the paper to point forecasts and have not discussed higher moments. Thus, it is important to emphasize that any published forecasts should come with measures of forecast uncertainty. Taking the unconditional forecast as a baseline, it would probably also be a good thing to publish one or more alternative scenarios based on conditioning assumptions regarding policy rates or other variables of particular interest, such as oil prices.

We now take up some objections to this proposal. It might be asserted that the CB generally has very little in the way of better information than the public. Thus, the analysis of the paper may not be very important. We agree that CBs in general are doing very well on transparency and there may not be huge gains left to realize. While this may be true on average or most of the time, there may come particular events when the CB and public views diverge. In such cases, having in place a good communication policy may be of considerable value.

Further, it is worth emphasizing that our argument need not rest on the CB having a *better* assessment of the economy than the public. Whether or not the CB forecast is better than that of the public, the forecast will be of value to the public since it will be the basis of policy. In addition, a clear statement of the bank's views will allow the public to monitor better the performance of the CB.

Another argument against our main conclusions is that using conditional fore-

[23] For example, what was the public to make of constant policy forecasts with inflation always returning to target combined with regular, serially correlated policy changes?

casts based on the market path attains most of the benefits of our proposal. We think this is not the case. While the market path has the benefit of probably constituting a modest policy intervention, other than this technical benefit, the market-based conditional is no better than any other. Perhaps the easiest way to see this is to consider the case in which the public's view of the future of policy is different from the CBs. While this may be the case in which a good inflation report is most needed, a forecast based on the public's own view of policy cannot reveal the CB's view.

Indeed, while the market rate assumption has no more information about policy than any other (modest) conditional, it is potentially more confusing. The public is not likely to confuse the constant rate assumption with the CBs view of where policy is headed. In the case of the market-based assumption, the public may naturally be inclined to wonder if the CB is ratifying the view in the market path.

The tension here is reflected in the different stances that the BOE and Riksbank take toward the market assumption. The Bank of England acknowledges that forward rates embody the market's view of future policy decisions and labels the projections—at least of inflation and output—the "best collective judgement" of the Monetary Policy Committee [Bank of England (2005)]. Taken literally, this seems to require, however, that the market's view of where policy is heading always reflects the best judgement of the Committee—even before the report is published. If this were not so, the the conditonal forecast should not reflect the committees view of the likely path of the economy. Of course, if the Committee's view of policy is already fully reflected in the market—even before publication of the report—then it is not clear what role the report is to play in transparency. In contrast, the Riksbank does not explicitly link forward rates to expected policy and its *Inflation Report* states that the conditioning path "should not be interpreted as the monetary policy assumption the Executive Board considers most probable" [Sveriges Riksbank (2005), p. 5]. If not, then what path is considered most probable? In our view the market-based conditional forecast answers few of our objections, but raises new questions.

The next objection to our proposal is that it is infeasible: unconditional forecasting is simply demanding. Unconditional forecasting is less demanding than conditional forecasting, in our view. Unconditional forecasting does not require making controversial identifying assumptions in order to estimate the impulse response to a policy shock, for example.

A different version of this objection is that at some central banks, such as the Federal Reserve, the sheer number of decisionmakers (19 on the FOMC) renders reaching a consensus on a forecast infeasible. This is a very serious issue and we do not take it up in this paper. At the inflation targeting banks that are currently publishing a conditional forecast, forming an unconditional forecast should not be more demanding than forming a conditional.

Finally, we return to the point that releasing a forecast of the policy rate would somehow damage the policymaking process. For example, the CB might later find it difficult to diverge from the point forecast, even when it is appropriate to do so. Alternatively, deviation from the point forecast might be seen as failure and, hence, diminish credibility. More generally, markets might simply overreact to rate forecasts.

We believe that if the rate forecast comes with an appropriate uncertainty measure, these problems may be minimized. Publishing alternative scenarios that show that policy will differ depending on how the economy evolves could also help. Banks such as the RBNZ and Bank of Norway are providing some empirical evidence on these claims. We have seen no measurable costs so far, but such costs might be hard to observe or arise only in especially challenging times. In the end, however, we cannot solidly reject the traditional view against central bankers discussing interest rates.

6 Conclusions

This paper argues that the most important aspect of central bank transparency is communication about the future path of the policy rate. Despite great progress on

transparency, communication on this topic remains shrouded in an obscure code. Central banks can enhance transparency by providing, in decreasing order of importance: i) an unconditional forecast of the policy interest rate; ii) an unconditional forecast of goal variables; iii) conditional forecasts of rates and goal variables. Of course, this is advice about increasing transparency, but does not establish that doing so would be good.

Of course, the first of our suggestions rejects one of the strongest taboos regarding central banker communication: never talk about the future of rates. Perhaps this taboo is well founded. As the experiences in New Zealand and Norway demonstrate, however, disaster need not follow discussing the future of policy in a relatively direct way. At the very least, it is time to thoroughly review the basis of this taboo.

Appendix

To be provided

References

Archer, David (2004), Communication with the Public, manuscript, Reserve Bank of New Zealand.

Bank of England (2005), *Inflation Report*, November.

Bank of England (2005), *Inflation Report*, May.

Bank of England (2004), *Inflation Report,* August.

Bernanke, Ben, Laubach, Thomas, Mishkin, Frederic, and Posen, Adam, *Inflation Targeting,* 1999, Princeton: Princeton University Press.

Eichenbaum, Martin., Evans, Charles., (1995). Some empirical evidence on the effects of shocks to monetary policy on exchange rates. *Quarterly Journal of Economics*, v.110, pp. 975–1010.

European Central Bank (2005), *Monthly Bulletin*, December.

Faust, Jon (2005), Is Applied Monetary Policy Analysis Hard? Manuscript, Federal Reserve Board.

Faust, Jon and Dale Henderson (2004), Is inflation targeting best practice monetary policy?, Federal Reserve Bank of St. Louis *Review*, v.86, n.4, pp. 117–144.

Faust, Jon and John H. Rogers (2003), Monetary Policy's Role in Exchange Rate Behavior, *Journal of Monetary Economics,* v.50, n.7 , October 2003, pp. 1403–1424.

Faust, Jon; Rogers, John H.; Wang, Shing-Yi; Wright, Jonathan H. (2003), The High Frequency Response of Exchange Rates and Interest Rates to Macroeconomic Announcements, manuscript.

Faust, Jon, Rogers, John H., Wright, Jonathan H., (2005) News and Noise in G-7 GDP Announcements, *Journal of Money, Credit, and Banking*, forthcoming.

Faust, Jon, Eric T. Swanson, and Jonathan H. Wright (2004) Do Federal Reserve Policy Surprises Reveal Superior Information about the Economy?, *Contributions to Macroeconomics,* v. 4, n.1, Article 10.

Faust, Jon and Jonathan H. Wright (2005), Econometric Analysis of Conditional Forecasts with an Application to Central Banking, in preparation.

Goodhart, Charles (2005), The Monetary Policy Committee's Reaction Function: An Exercise in Estimation, manuscript.

Greenspan, Alan (2002), Chairman's Remarks, Federal Reserve Bank of St. Louis *Review*, v.84, n.4, pp. 928–56.

Issing, Otmar (2005), Communication, Transparency, Accountability: Monetary Policy in the Twenty-First Century, Federal Reserve Bank of St. Louis *Review*, v.87, n.2, Part 1, pp. 65–83.

Leeper, Eric M. and Zha, Tao, Modest Policy Interventions, *Journal of Monetary Economics,* vol. 50, iss. 8, 2003, pp.1673–1700.

Leeper, Eric M. (2003), An 'Inflation Reports' Report, Sveriges Riksbank *Economic Review*, n.3., pp. 94–118.

Norges Bank (2005), *Inflation Report with Monetary Policy Assessments*, 3, November.

Norges Bank (2005), *Inflation Report*, March.

Pagan, Adrian, *Report on modeling and forecasting at the Bank of England*, London: Bank of England, 2003.

Reserve Bank of New Zealand (2005), *Monetary Policy Statement*, December.

Reserve Bank of New Zealand (2005), *Monetary Policy Statement*, March.

Romer, Christina and David Romer (2000). Federal Reserve Information and the Behavior of Interest Rates, *American Economic Review,* 90, pp. 429–57.

Sims, Christopher A. (2002), The Role of Models and Probabilities in the Monetary Policy Process, *Brookings Papers on Economic Activity,* iss. 2, pp.1–40.

Sims, Christopher A., Macroeconomics and Reality *Econometrica,* vol. 48, iss. 1, 1980, pp.1–48.

Svensson, Lars E.O. (2001), Independent Review of the Operation of Monetary Policy in New Zealand: Report to the Minister of Finance, February 2001.

Svensson, Lars E.O. (2004), Monetary Policy with Judgment: Forecast Targeting, manuscript, Princeton University.

Svensson, Lars E.O. and Noah Williams (2005), Monetary Policy with Model Uncertainty: Distribution Forecast Targeting, manuscript, Princeton University, May.

Sveriges Riksbank (2005), *Inflation Report 2005:4*, December.

Sveriges Riksbank (2005), *Inflation Report 2005:1*, March.

Sveriges Riksbank (2003), *Inflation Report 2003:4*, December.

Waggoner, Daniel F. and Zha, Tao (1999) Conditional Forecasts in Dynamic Multivariate Models, *Review of Economics and Statistics*, v. 81, n.4, pp. 639–651.

Woodford, Michael (2005), Central-Bank Communication and policy Effectiveness, manuscipt, Columbia U.

Woodford, Michael (2004), Inflation Targeting and Optimal Monetary Policy, Federal Reserve Bank of St. Louis *Review*, v.86, n.4. pp. 15–42.

Table 1. Impulse response of interest rate and output and implied Γ-response of output in a standard monetary VAR

hor:	0	1	2	3	4	5
i	1.00	1.52	1.34	1.25	0.99	0.60
100 b_h	0.00	-0.25	-0.58	-1.19	-1.79	-2.35
100 γ_h	0.00	-0.25	-0.20	-0.55	-0.37	-0.55

Notes: The estimates come from Faust-Rogers (2003) replication of an open economy VAR of Eichenbaum and Evans (1995). These are for a 7-variable model of the U.S. and U.K. and represent the response of U.S. output growth to a policy shock that initially raises the policy rate by 1.

Table 2a. Γ-Responses for inflation implied by pairs of BOE Forecasts

hor:	0	1	2	3	4	5
1998:1	0.00	0.00	-2.37	11.75	-2.77	-215.35
1998:2	0.00	0.00	-0.50	3.14	-9.85	21.43
1998:3	0.00	0.00	0.00	0.00	0.00	0.00
1998:4	-0.54	-3.81	-25.61	-187.24	.	.
1999:1	16.57	.	.	.	.	.
1999:2	0.00	0.00	-0.23	0.52	-0.84	1.15
1999:3	0.39	-0.66	0.27	1.53	-1.27	3.39
1999:4	0.00	-0.29	-0.08	-1.04	-1.63	-7.39
2000:1	0.00	0.70	4.42	32.44	230.24	.
2001:1	0.00	0.00	1.32	16.78	216.94	.
2001:2	0.00	0.00	0.00	0.00	1.20	6.73
2001:3	0.00	0.00	0.00	0.00	0.00	-0.56
2001:4	0.00	0.00	0.00	0.00	-0.23	0.69
2002:2	0.00	0.00	-1.32	8.62	-31.29	73.07
2003:1	0.26	-0.34	0.45	-0.18	-0.25	-0.23
2003:2	0.00	0.00	0.00	1.36	-18.01	203.36
2003:4	0.00	0.00	1.30	15.29	208.80	.
2004:1	0.00	0.00	0.00	0.63	3.32	15.69
2004:2	0.00	0.00	0.00	1.26	12.66	119.47
2004:3	0.00	0.00	0.00	-0.18	0.35	-0.88
2004:4	0.00	0.00	0.00	0.00	0.14	-0.22
ave.	0.79	-65.66	.	.	.	.
min.	-0.54	.	-25.61	.	.	.
max.	16.57	0.70	.	32.44	.	.

Notes: See the notes at end of Table 2b.

Table 2b. Γ-Responses for output growth implied by pairs of BOE Forecasts

hor:	0	1	2	3	4	5
1998:1	21.19	-133.44	204.58	.	.	.
1998:2	3.01	-14.28	35.79	-69.56	120.46	-192.33
1998:3	0.47	-1.10	0.52	-0.99	1.60	-0.34
1998:4	-1.64	-7.16	-54.84	-393.42	.	.
1999:1	-8.40	683.80	.	.	.	.
1999:2	0.00	-0.24	0.29	-0.10	0.50	-1.01
1999:3	0.00	0.00	1.16	-1.21	3.75	-0.58
1999:4	0.00	0.00	0.58	1.04	3.76	9.65
2000:1	0.00	0.69	6.47	44.74	320.58	.
2001:1	0.00	2.63	36.11	468.87	.	.
2001:2	0.00	1.20	7.91	49.16	298.21	.
2001:3	0.00	-0.56	0.53	-0.67	0.32	-0.21
2001:4	0.00	-0.23	0.23	-0.23	-0.19	-1.11
2002:2	0.00	-0.66	0.72	-0.37	-2.42	12.12
2003:1	-0.26	0.08	-0.11	-0.02	0.08	0.16
2003:2	-1.36	15.29	-172.84	.	.	.
2003:4	0.00	1.30	19.20	259.86	.	.
2004:1	0.00	2.49	10.55	49.26	232.73	.
2004:2	1.24	16.11	146.79	.	.	.
2004:3	0.00	-0.35	0.87	-1.89	3.88	-7.29
2004:4	0.00	0.14	-0.22	0.13	-0.07	0.32
ave.	0.68	26.94	.	.	.	.
min.	-8.40	-133.44	.	-393.42	.	.
max.	21.19	683.80	204.58	.	.	.

Notes: These are the Γ-responses as described in the text computed from pairs of BOE forecasts where the forecasts are taken from the IR with the date noted in the relevant row. Where there is only a dot, the value in question was greater than 999. Rows for a few IRs are missing; this occurs when the pair of interest rate paths does not differ sufficiently to meet the invertibility condition.

Table 3a. Bank of England estimated Γ-responses and forecast efficiency tests

h:	0	1	2	J-test
		BOE, inflation, system		
γ	-26.06	.	.	.
se	(64.93)	.	.	.
γ	-0.86	-0.99	.	0.77
se	(1.24)	(11.04)	.	(0.68)
γ	-1.30	3.01	7.83	2.47
se	(1.40)	(4.51)	(17.89)	(0.78)
		BOE, inflation, single		
γ	-21.38	.	.	0.01
se	(113.94)	.	.	(0.94)
γ	1.24	-40.94	.	0.01
se	(7.17)	(148.96)	.	(0.92)
γ	-3.48	4.42	51.98	0.04
se	(21.55)	(24.03)	(188.42)	(0.84)
		BOE, GDP, system		
γ	75.28	.	.	.
se	(158.32)	.	.	.
γ	-1.61	-6.44	.	3.50
se	(2.36)	(22.29)	.	(0.17)
γ	-1.69	-0.79	19.98	4.93
se	(0.89)	(2.55)	(18.69)	(0.42)
		BOE, GDP, single		
γ	-42.88	.	.	3.05
se	(19.55)	.	.	(0.08)
γ	-1.00	2.84	.	1.72
se	(4.77)	(58.17)	.	(0.19)
γ	2.80	-4.51	-30.64	3.04
se	(3.50)	(4.95)	(19.69)	(0.08)

See notes at end of table.

Table 3b. Riksbank: estimated Γ-responses and forecast efficiency tests

h:	0	1	2	J-test
		Riks, inflation, system		
γ	6.99	.	.	.
se	(4.57)	.	.	.
γ	1.56	-8.04	.	0.49
se	(6.24)	(31.03)	.	(0.78)
γ	-1.76	8.70	-27.66	1.91
se	(2.05)	(5.07)	(11.08)	(0.86)
		Riks, inflation, single		
γ	9.78	.	.	0.34
se	(8.62)	.	.	(0.56)
γ	1.11	-5.52	.	0.61
se	(7.60)	(40.20)	.	(0.43)
γ	-4.96	6.36	3.34	0.22
se	(12.05)	(19.03)	(28.88)	(0.64)
		Riks, GDP, system		
γ	-8.81	.	.	.
se	(7.88)	.	.	.
γ	-8.40	40.60	.	0.20
se	(26.84)	(136.78)	.	(0.91)
γ	-4.01	21.16	-92.79	3.74
se	(1.35)	(6.79)	(30.64)	(0.59)
		Riks, GDP, single		
γ	-44.96	.	.	2.46
se	(26.10)	.	.	(0.12)
γ	-8.83	16.21	.	0.85
se	(17.18)	(85.76)	.	(0.36)
γ	-4.48	4.43	-1.27	1.20
se	(12.16)	(20.74)	(28.16)	(0.27)

See notes at end of table.

Table 3c. Fed: estimated Γ-responses and forecast efficiency tests

h:	0	1	2	J-test
		Fed, inflation, system		
γ	0.23	.	.	.
se	(1.30)	.	.	.
γ	0.10	-1.90	.	0.03
se	(0.33)	(1.07)	.	(0.98)
γ	0.34	-0.89	-1.97	5.44
se	(0.33)	(0.58)	(1.15)	(0.36)
		Fed, inflation, single		
γ	0.17	.	.	0.35
se	(1.29)	.	.	(0.55)
γ	-0.14	-1.01	.	2.67
se	(0.76)	(2.37)	.	(0.10)
γ	0.32	-0.36	-1.56	3.12
se	(1.13)	(2.38)	(2.23)	(0.08)
		Fed, GDP, system		
γ	-2.72	.	.	.
se	(1.96)	.	.	.
γ	0.43	-1.84	.	2.76
se	(0.77)	(2.37)	.	(0.25)
γ	0.90	0.54	2.96	8.89
se	(0.67)	(1.27)	(2.94)	(0.11)
		Fed, GDP, single		
γ	-1.96	.	.	0.62
se	(2.45)	.	.	(0.43)
γ	2.60	-6.05	.	1.56
se	(1.43)	(5.45)	.	(0.21)
γ	1.31	1.83	-3.77	1.09
se	(10.74)	(20.99)	(17.75)	(0.30)

See notes next page.

Notes: Each row of the table presents the results of a separate 2-step GMM estimation of a system of the form (25). The rows labelled γ give the point estimate and the rows labelled *se* give the standard error for the estimate just above. The J-test column gives the value of the J test statistic and associated p-value in parenthesis. In the *system* blocks, the model is a joint treatment of the equations for horizons zero through h with the γs the same across equations. In the *single* blocks, the estimation is for the equation for horizon h. Instruments for the various models are noted in the text.