FORECASTS AND INFLATION REPORTS: AN EVALUATION

JON FAUST AND ERIC M. LEEPER

Abstract. This paper argues that the most straightforward way central banks can enhance transparency is to provide the following, in 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. The paper shows theoretically that it is very difficult to infer unconditional forecasts—the objects that reflect central banks' full information set—from conditional ones—the objects banks rountinely publish. It evaluates the information content of actual central bank forecasts and extracts practical implications from the theoretical and empirical results.

1. Introduction

For most of the history of central banking, the conventional wisdom on transparency was captured by the adage "silence is golden." Those days are over. A general consensus has emerged that central banks should fully reveal policy-relevant information, unless there is clear justification for maintaining secrecy. The inflation targeting central banks have in many cases led the way in this regard—transparency is at the core of the inflation targeting framework. Most inflation targeters publish regular inflation reports with extensive examination of economic issues driving monetary policy. One important element of these reports is some sort of information regarding forecasts of important variables.

Few would argue, however, with the view that there are limits to how far central banks can beneficially push transparency. As Chairman Greenspan (2002, pp. 5–6) of the Fed noted, "Openness is an obligation of a central bank in a free and democratic society," but also, "The undeniable, though regrettable, fact is that the most effective policymaking is done outside the immediate glare of the press."[1]

Date: June 6, 2005. Federal Reserve Board; faustj@frb.gov (http://e105.org/faustj) and Department of Economics, Indiana University; eleeper@indiana.edu http://mypage.iu.edu/~eleeper/).
We thank Jonathan Wright for constructive conversations. The paper was prepared for the Sveriges Riksbank conference "Inflation Targeting: Implementation, Communication and Effectiveness," June 11-12, 2005. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or other members of its staff.

Despite the remarkable progress on transparency, it makes sense to ask whether central banks have reached the efficient frontier in the tradeoff between transparency and other considerations. This paper takes up one aspect of this question: Can forecast information be used more effectively in communicating the central bank's views of policy and the economy? We will argue strongly that it can.

The basic argument goes like this. Since the central bank's main public policy mandate is to set monetary policy, the primary focus of transparency should be to reveal the bank's views of the future course of policy. Additionally, transparency dictates that the central bank discuss the implications of policy for the economy. If the central bank uses an interest rate as the primary instrument, straightforward application of this idea would have the central bank give its best guess as to the likely course of interest rates. This brings us to what may be the last remaining area in which central banks cling to the view the silence is golden.

It is natural to couch this discussion in terms of forecasts. Inflation reports generally contain one or more *conditional forecasts*: forecasts conditioned on some exogenously specified path of the policy interest rate such as a constant policy rate. If, as we argue, the primary focus of transparency should be on revealing the central bank's best guess of the likely course of interest rates and the goal variables of policy, it would be more natural to report the central bank's *unconditional forecast*—the central bank's best view of where the economy is headed conditioned only on information available to it at the time of the forecast.

After more fully supporting our view of the goals of transparency, much of the paper analyzes the relative information content of conditional and unconditional forecasts. We present theoretical results for a best-case scenario for conditional forecasts—a simple model populated by agents with a great deal of sophisticated knowledge about the nature of the economy and policy. Even in this case, conditional forecasts cannot convey important central bank information. Of course, reality is much messier: the *model* is extremely complicated and not all agents (perhaps no agents) are as knowledgeable and sophisticated as those in our models.

In a realistic context, the limitations of conditional forecasts are much more severe. Even among professional economists there are many outstanding issues regarding the proper analysis of conditional forecasts in practice. For example, when assessing the accuracy of central bank forecasts econometrically, standard practice is to treat the forecasts as *unconditional.* While central bank forecasts have generally been found to be pretty good when viewed as unconditional forecasts, these papers generally note that this is not necessarily a good thing: good prediction based on a counterfactual

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

policy assumption could, in principle, be evidence of a flawed forecast process. This paper and a related one by Faust and Wright (2005) derive and apply some new results about the empirical analysis of conditional forecasts. We find room for improvement.

While these new methods and the associated empirical results have some implications for improving central bank forecasts, the main focus of this paper is on transparency. What can the public learn from forecasts? From this standpoint, the complexity of deriving and applying the methods drives home the point that conditional forecasts are difficult to formulate, understand, and analyze. They are not, we argue, the most straightforward way to communicate. In the penultimate section, we discuss a number of ways to improve both the forecasts and their contribution to achieving the goal of transparency.

Our main conclusion: the most straightforward way to enhance transparency is for the central bank to provide the following, in order of importance: i) an unconditional forecast of the policy interest rate; ii) an unconditional forecast of goal variables such as inflation and activity; and iii) conditional forecasts of rates and goal variables.

While we illustrate our points technically, if difficult technique is required, perhaps one should seek a simpler way to communicate the ideas. In this spirit, we give a simple example. The policymaker is a bus driver and the public is the passengers. Full transparency in this case dictates that the bus driver share her private information about when the bus will arrive at the destination. The typical passenger has a rough idea of the travel time, but the bus driver, who is a regular on this route, has some private information. A passenger inquires about the expected travel time.

Rather than giving an unconditional forecast (the trip usually takes 48 hours), the bus driver reports two conditional forecasts for certain 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.

These are interesting facts, but how can the passenger use them to update his rough view of the likely travel time? If the passenger 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 back out the bus 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 estimate of arrival time.

In short, it is specialized knowledge that allows the policymaker to make sophisticated conditional forecasts. Without nearly equivalent specialized knowledge, it is difficult for the public to properly digest and benefit from announcements regarding these conditional forecasts.

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. We return to this point in the conclusion.

It might be best to stop the paper here—technical elaboration of these points may serve to obscure as much as illuminate. But we press on. Section 1 provides background on transparency and forecasts. Sections 2 and 3 discuss theoretical and empirical issues regarding the information content of forecasts, respectively. Section 4 provides some speculation about practical ways to improve transparency and Section 5 concludes.

2. Background on Transparency and Forecasts

2.1. **Some Principles of Central Bank Transparency.** We can codify the view of transparency in the introduction in terms of three principles of increasing specificity.

First, absent competing interests, central banks should strive to shed as much of their private information as is reasonably possible.

The basis of this principle is not that shedding private information is clearly welfare enhancing. As the transparency literature demonstrates, there are many reasons why in a distorted economy it might be welfare enhancing for the central bank deliberately to surprise the public.[2] But there is a strong consensus among academic economists and central bankers that in a democracy, this is an inappropriate thing for a central bank to do.

Most prominent among the competing arguments is the possibility that, due to some sociopolitical factors, certain types of transparency might distort the policy formation process. Some degree of secrecy might help the central bank resist the occasional political pressure for bad policy. We set this issue aside for now, returning to it only at the end.

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;[3] the central bank has a particular

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

[3] Central banks around the world differ in what other responsibilities they have beyond monetary policy making. These may be regulatory, participating in the payments system, roles as fiscal agent, etc. We do not mean to belittle these. Thus, this principle should probably begin "In the traditional monetary policy arena."

responsibility to share information about its primary mandate. More significantly, 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. For various reasons, central banks expend great resources evaluating the state of the economy and producing forecasts. Since these activities are costly and the information obtained has public good qualities, it might even be socially efficient for the central bank to play a unique information-generation or information-aggregation role.

Thus, along with monetary policymaking, the central bank could play a role like that of the National Weather Service in the United States in providing a forecast of important variables as a public good. For the purposes of this paper, this is a secondary or derivative responsibility in fulfilling the transparency principle stemming from the central bank's conduct of its core mandate.

Third, in pursuing transparency a central bank using an interest rate instrument also should be transparent about the future path of instrument.

While the first principle seems to have widespread support inside and outside central banks, as we move to the second, things get murkier, and by the third many central bankers are quite uncomfortable. This is presumably because the sociopolitical problems are perceived to be most acute where the instrument is concerned. For now we want to emphasize that the second two principles follow directly from the first.

An alternative view is that it is sufficient for the central bank to discuss only the goal variables of policy. This view can be summarized as follows. Suppose aggregate welfare can be approximated in the standard way as being quadratic in the output gap and the gap between actual and desired inflation. More generally, suppose that these two gaps (approximately) form a sufficient statistic for aggregate welfare. In this case, it might seem that the central bank could focus only on communication about this sufficient statistic.

To return to the bus driver example, it is important that the bus driver communicate about goals such as travel time and avoiding collisions. It is not nearly as important that she communicate about the manipulation of the instruments—how the steering wheel and gas pedal will be adjusted—to achieve the goals.

Similarly, discussion of the two gaps may be sufficient to communicate about *aggregate welfare*. Individual welfare, however, is strongly affected by the path of the instrument. For example, changes in nominal interest rates have tremendous power to redistribute wealth. Perhaps appropriately, we largely neglect these in aggregate welfare considerations. It is clear, however, that a person in a very secure job, but facing interest-sensitive decisions may care more about the short-run path of interest rates than about the short-run path of the gaps. A simple thought experiment—repeated regularly in the lives of many central bankers—may be instructive. A person says to the central banker something like, "I am about to lock in a fixed-rate mortgage. The consensus seems to be that the central bank won't lower rates for a while. What do you think?" If, in my role as a central banker, I have reason to believe that the consensus is wrong, should I share this information? Under the three principles outlined here, the central bank's characterization of where rates are headed is in the public domain already. The central banker could review this information, noting where it seems to differ from what is taken to be the consensus. In contrast, under the situation prevailing at most central banks, the appropriate answer is "I can't comment." Perhaps she adds, "If the policy rate were to remain constant, output growth would do..., and inflation would...."

Absent competing considerations, if the central bank has better information about the course of interest rates than the public, the transparency principles imply that this information should be shared. This paper discusses how effectively banks currently share this information and suggests how they might do so more completely. Our goal is to fill out the benefit side of the cost-benefit analysis that underlies a choice of optimal transparency.

2.2. **Forecasts and the Central Bank Communication Problem.** Especially since many central banks started publishing regular forecasts of the economy, the role of forecasts in optimal policymaking and in transparency has appropriately received great attention. We distill some key results from this literature here. While we are not sure these points have been collected this way, they are present in many places (e.g., Sims, 1980; Bernanke, et al. 1999; Leeper and Zha 2003; Svensson 2004, Woodford 2004)

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

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

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

Without further structure, the policy communication process is very difficult. To communicate successfully about the future course of policy, the central bank might need to write an extended discussion of current information and its understanding of the map between information and the policy choice.

The goal in designing a communication policy is to reduce or condense the characterization of the current information and the policy mapping to something that can effectively be communicated. Some conventional assumptions simplify the problem considerably.

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

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

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 if the history is Y_{t-1} and the policy rate is $i_t = i$. Note that we are not necessarily assuming that L is a true social welfare function, so this expression covers a wide range of cases. Of course, the density $\xi(Y|i, Y_{t-1})$ can be called a forecast density; thus, equation (1) emphasizes the tautology (by the meaning of expected) that, period-by-period expected loss minimization can be written in terms of forecasts.

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

$$L(Y_t) \approx l(Z_t),$$

where Z_t is a function of Y_t and of much lower dimension. In the familiar case, this function picks out the paths of the two gaps. Now we can write,

$$i_t \approx \arg\max_i E\left[l(Z)|Y_{t-1}\right] = \int l(Z)\xi(Z|Y_{t-1})dZ.$$

The communication problem is now much simpler. Now the central bank only needs to communicate its preferences and a forecast for the low dimensional variable Z. This forecast itself is conditioned on the full information set, but the forecast of the Z's alone is sufficient for explaining policy.

2.3. **Conditional and Unconditional Forecasts.** The remaining communication problem is not easy; it is somewhat simplified relative to the original. The central bank is still faced with the daunting task of reporting a multi-dimensional conditional forecast density, ϕ, for a continuum of values of the multi-dimensional conditioning variable, i_t.

We will call forecasts conditional on central bank information at the time of the forecast and a future path for the policy rate *conditional* and contrasting that with a forecast conditional only on central bank information, *unconditional.* As we understand it, central banks have generally not reported an unconditional forecast in their inflation reports.

Most central banks (with the exception of the BOE) have generally reported one conditional forecast, but an increasing number of central banks are publishing multiple sets of conditional forecasts. In addition to the "main scenario," which usually conditions on a constant interest rate path, some banks consider an alternative path for i drawn from futures markets.[4]

For our analysis, it is important that the banks present one or two conditional forecasts—the particulars of the conditioning path are not crucial.

Because we focus on the distinction between conditional and unconditional forecasts, one point about conditioning on the futures path is important. It might seem like conditioning on the futures path would be close to the unconditional forecast. For example, suppose that the forecasting model is linear. If one conditions on the unconditional expectation for some variable, the resulting *conditional* expectation for all the variables is the *unconditional expectation.* Thus, if the forecasting model of the central bank is linear and the path of i implied by market rates corresponds to the central bank's unconditional forecast, the resulting forecast will be close to the unconditional forecast.

First, the market-based conditional will only be the unconditional in the case that the market-implied path of the policy rate is equal to the central bank's unconditional expectation for the policy rate. Thus, a market-based conditional forecast corresponds to the unconditional *only when there is no need to communicate about the future course of policy.* When there is a communication problem, the two will not correspond.

We will see this issue in many guises: the conditional forecasts may intuitively seem informative, but careful analysis shows that this intuition requires that the public interpreting the forecast has equivalent, or nearly equivalent, information to the central bank. Of course, in this case, the forecasts merely repeat information; they do not expand the information of the public. We should note that *repeating* information may be a very good thing. It puts it in the public record and may have other useful purposes. Our focus, however, is strictly on whether the forecasts in inflation reports help shed private information of the central bank.

[4] Banks now conditioning on market interest rates include at least England, Norway, and Sweden. The first two treat forecasts conditioned on market rates as their central projections, while the Sveriges Riksbank treats the constant interest rate forecast as its main scenario.

A second point is worth noting. The first moment of the market-based conditional and unconditional will correspond in the case just noted, but higher order moments will not. The conditional forecast takes the policy path as constant, while it will be uncertain in the unconditional.

For simplicity we are going to set aside the complicated analysis of higher order moments of conditional and unconditional forecasts. Other papers have dealt with some important theoretical and empirical issues in this regard (Svensson, 2004,2005; Sims 2002).

We focus attention on the simple question raised in the introduction regarding the bus driver. When do conditional forecasts reveal the information that would be reflected in the unconditional forecast? This is both the simplest and arguably the most important issue to address: If the central bank has a better idea than the public about the expected path of the policy rate and the economy, *absent other considerations* it should share it. Higher order moment information is also useful, but as we know, this is likely to have smaller welfare implications.

Two complimentary technical views of our approach may help motivate our analytical perspective. The first views the publication of forecasts as a signalling problem. The central bank maps all its information into a signal in the form of the announced forecast paths. Can this mapping be inverted by the public to recover all the information? If the mapping were linear, there would be a *spanning* or full-rank condition governing the invertibility of the mapping. In particular, the degree of private central bank information would be measured by the rank of the space spanned by the private information, and if the central bank announcement spans this space, then the announcement is fully revealing.

Alternatively, we can view this as an econometric identification problem. Suppose we give an econometrician the public information set plus the set of the forecasts announced by the central bank over several successive announcements. What aspects of the central bank's private information would be econometrically identified? These two perspectives are complementary, and, of course, a rank condition is at the core of identification.

By focusing on the first moment of forecasts, we run the risk of deflecting attention from one of Sims's (2002) key points: central banks should focus more on complete probability models. We agree with this point, but accept the risk in order to forward our main point, which we view as equally important. The next two sections examine theoretical and empirical aspects of our main point about the information content of conditional and unconditional forecasts, respectively.

3. Conditional and Unconditional Forecasts in Theory

3.1. **An Illustrative Model.** This section lays out a stylized model in which there is a role for inflation reports to inform the public. The model's unconditional and conditional forecasts of macro variables—possibly including the policy interest rate—potentially contain different information. It is pretty easy to generate a case in which the conditional forecasts contain essentially no useful information about the unconditional.[5] We focus on a very simple setup that offers a best-case environment for inflation reports as they are currently produced.

3.1.1. *The model and its equilibrium.* The economy is described by a simple Phillips curve and the central bank optimizes a loss function that is quadratic in an output gap and inflation. The transparency issue arises because the central bank receives a private signal about the state of the economy. In simple models, there is always a solution to such communication problems: tell people your signal.[6] We set aside this solution in order to illustrate some key points.

We model the central bank's information advantage as arising from resources it devotes to data collection and analysis and to econometric modeling of the economy. The central bank continuously optimizes, reacting to the signals it receives about the state of the economy. All dynamics are compressed into a single period, which we take to be on the order of six months to a year. The central bank correctly perceives that it faces a tradeoff in achieving its twin objectives of inflation and output stabilization. Its choice of nominal interest rate, i, gives the central bank a direct impact on inflation, π, and the output gap, x.

The Phillips curve is given by

$$x_t = -ai_t + b(\pi_t - \pi_t^e) + u_{xt} + w_{xt} \tag{2}$$

where x_t and π_t are publicly observed measures of the output gap and inflation at t, respectively. Further, π_t^e is the public's expectation of inflation formed at the end of $t-1$ and u and w are shocks that are not directly observed by anyone. Inflation is given by,

$$\pi_t = -ci_t + u_{\pi t} + w_{\pi t} \tag{3}$$

[5] For example, if the central bank's model has constantly shifting parameters known only to the central bank.

[6] To be clear, we are not studying cases in which the central bank has any tendency to lie to the public.

where u and w are again shocks. For later reference, it is useful to define the vector $Z_t = (x_t, \pi_t)'$ and solve the resulting system for Z:

$$Z_t = Ci_t + B(u_t + w_t) \tag{4}$$

where

$$C = \begin{bmatrix} -(a+bc) \\ -c \end{bmatrix}, \qquad B = \begin{bmatrix} 1 & b \\ 0 & 1 \end{bmatrix},$$

and u and w are (2×1) vectors stacking the two u and w shocks, respectively. In our simple static economy, C gives the impulse response of the Z variables to a monetary policy shock.

The central bank chooses the policy rate i_t at the beginning of time t to minimize what we stipulate to be the social loss function given by:

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

As we shall see, the central bank makes its choices after observing signals about the of the shocks, so the information set associated with the operator E_t includes these signals. All parameters in the model are positive. The shocks are normalized to have mean zero and are *i.i.d.*

It is the interpretation of the shocks and when they enter the information sets of the public and central bank that is central to our discussion. The pair of shocks u_x and w_x raise x_t, the public measure of the gap, but it is only $x_t - u_{xt}$ that enters the loss function; in this sense, movements in x due to u_x are efficient and should not be offset; movements due to u_π lower welfare. The analogous story holds for π, u_π and w_π.

The key idea we hope to capture is that public measures of inflation and the gap fluctuate for many reasons and the central bank appraises many diffuse sources of information in deciding how to respond to any given movement. We model this phenomenon in a reduced-form manner by assuming that the central bank receives an unbiased private signal about each of the four shocks: $\hat{v}_t = (\hat{u}_{xt}, \hat{w}_{xt}, \hat{u}_{\pi t}, \hat{w}_{\pi t})$.

The timing of events is as follows. At the end of $t-1$ the public forms it expectation π_t^e. Period t then has four events. At the beginning of t, the central bank's signal, $\hat{v}_t$, arrives. Next the central bank sets i_t optimally, taking π_t^e as given; then the shocks, v_t, are realized and the private sector chooses x_t and π_t; finally, the public forms its expectation for $t+1$ and the economy begins again.

Now we overlay this structure with the communication issue. 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 conveyed by various sorts of forecasts. One might wonder why the public would bother to read the inflation report: the information cannot affect outcomes as we have described the model. This feature keeps things simple. We are mainly interested in what information can be conveyed in forecasts and it is not essential that we model effects that might emerge. We are free, however, to follow the discussion above in imagining that agents make additional decisions after reading the inflation report, which through distributional effects, change individual, but not social welfare.

Given the linear quadratic structure, optimal policy will be of the form,

$$i_t = \alpha_1 \hat{u}_{\pi t} + \alpha_2 \hat{w}_{\pi t} + \alpha_3 \hat{w}_{xt} + \alpha_4 \pi_t^e. \tag{5}$$

Policy tries to smooth the impacts of the w's; it ignores u_{xt} because output fluctuations from this source are efficient; it responds to $u_{\pi t}$ because by (2), movements in π due to either shocks lead to welfare losses through x.

The first-order conditions for the central bank imply

$$E_t(\pi_t - u_{\pi t}) = \frac{\lambda(a+bc)}{c} E_t(x_t - u_{xt}). \tag{6}$$

Because the private sector does not observe the central bank's signals, even under optimal policy there is a component of policy behavior that is random to the private sector. Given that the shocks are mean zero, this expression implies that

$$\pi_t^e = 0.$$

Neither the central bank nor the public at the end of $t-1$ have any information that would suggest an efficient deviation from zero inflation.

Letting $Y_t = (x_t, \pi_t, i_t)'$ and distinguishing between the realizations of the shocks, $v_t = (u_{xt}, w_{xt}, u_{\pi t}, w_{\pi t})$, and their signal, $\hat{v}_t$, straightforward algebra yields the equilibrium:

$$Y_t = A_0 v_t + A_1 \hat{v}_t, \tag{7}$$

where

$$\begin{aligned} A_0 &= \begin{bmatrix} 1 & 1 & b & b \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \\ A_1 &= \frac{1}{\Delta} \begin{bmatrix} 0 & -\lambda(a+bc)^2 & -\lambda b(a+bc)^2 & (a+bc)\left[-\lambda b(a+bc)+c\right] \\ 0 & -\lambda c(a+bc) & -\lambda bc(a+bc) & -\lambda bc(a+bc)+c^2 \\ 0 & \lambda(a+bc) & \lambda b(a+bc) & \lambda b(a+bc)-c \end{bmatrix}, \end{aligned}$$

with $\Delta = \lambda(a+bc)^2 - c^2$. The third row of (7) reports the α_i's in the optimal rule, (5), after imposing that equilibrium implies $\pi_t^e = 0$.

3.1.2. *The information content of conditional forecasts.* After receiving the signal, $\hat{v}_t$, the central bank computes both unconditional and conditional forecasts. The conditional forecast is conditioned on some policy assumption, say, $i_t = i_t^c$. That forecast, which is given by inserting $i_t = i_t^c$ in (2) and (3) and placing hats on all the shocks, is

$$Z_t^c = Ci_t^c + B(\hat{u} + \hat{w}). \tag{8}$$

Given that (4) must hold in equilibrium, even without solving for the unconditional forecast, we know that (8) must hold for any forecast including the one we will call unconditional; thus,

$$Z_t^c - Z_t^u = C(i_t^c - i_t^u) \tag{9}$$

The forecasts differ by the impulse response of a policy shock times the "policy shock" implied by the difference between the conditional and unconditional forecasts.

Note that forming the conditional forecast requires no use of the model solution—it thereby cannot reveal much about policy behavior, whether optimal or not. Clearly, if the public knows the model, this conditional forecast is equivalent to announcing the pair $(\hat{u}_{\pi t} + \hat{w}_{\pi t})$ and $(\hat{u}_{xt} + \hat{w}_{xt})$, whereas the public would like to know the values of the four shocks separately.

More generally, this exercise illustrates several important points that we will generalize below. First, in conditional forecasting the central bank suspends the actual interest rate equation and, therefore, the restrictions imposed by optimal policy. This masks information about the nature of shocks and about the central bank's response. Such information could be of value to the public.

Second, the conditional forecast purges any information about the coefficients of the optimal policy rule—the α's in (5).

Third, although conditional forecasts do not reveal the central bank's private information, a pair of conditional forecasts does reveal something about the response of the economy to policy shocks. Specifically,

$$Z_t^{c_1} - Z_t^{c_2} = C(i^{c_1} - i^{c_2}). \tag{10}$$

So long as the central bank announces the conditioning assumption (which inflation reports do), two conditional forecasts immediately imply C. In this static economy this is the only non-zero term in the impulse response to a policy shock. In the dynamic example below, we show that the contemporaneous impulse response is revealed, but for higher order responses only something related to the conventional impulse response is revealed. In any case, the BOE has long produced two conditional forecasts in its inflation report and we examine the C parameters implied by differencing below.

Without multiple conditional forecasts at a single point in time, the public cannot directly calculate C. But, as Faust and Wright (2005) show, a time series of such forecasts from successive inflation reports does allow one to estimate C. We report and extend some of Faust and Wright's empirical results below.

3.1.3. *The information content of an unconditional forecast.* As we noted above, what is known as the unconditional forecast in this context is the central bank's unbiased prediction of the economy given its information (including its knowledge of optimal policy). The unconditional central bank forecast, denoted by (x_t^u, π_t^u, i_t^u), is given by evaluating (7):

$$Y_t^u = A_1 \hat{v}_t.$$

If the most likely paths for the trio (x_t, π_t, i_t) are sufficient for individual welfare, the unconditional forecasts are all that are needed in an inflation report. In this case, private agents are made no better off by learning the precise decomposition of the central bank's information among the u's and w's.

It turns out that even in this very simple model, even the combination of conditional and unconditional forecasts alone cannot reveal all the central bank's private information. The reason is that optimal policy behavior imposes restrictions that reduce the rank of A_1 to 1. Conditional forecasts can contain information that is independent of unconditional forecasts because they embed counterfactual policy behavior, but in this model the additional information still does not permit agents to infer all of the central bank's information.

One might have the urge to derive conditions in which conditional and unconditional forecast paths are fully revealing. We do not want to push this little model that far. In practice, the knowledge we are talking about is quite diffuse and the conditions for inflation reports to be *fully revealing* from any simple model are probably not helpful. Rather, we are thinking about practical ways to give the public information that is more nearly sufficient. Our main point is that an unconditional forecast for the policy rate and the economic variables is the next natural step to consider in expanding central bank transparency.

3.2. **More General Theory.** Almost all the key points we wish to make are present in the simple model, but some subtleties emerge in the more general dynamic case. As noted above, very little in what we wish to emphasize depends on the particular nature of the optimizing behavior of the central bank or private agents. Thus to keep things a bit simple in the more general case, we merely posit that the central bank has a model, and do not derive the model explicitly.

The main goal is to derive the equivalent of (9) in the dynamic case, for this relation holds essentially all the information about how much information in the unconditional forecast can be inferred from the conditionals.

As before, the information set is Y_t, which now is a vector of explicitly measured variables. Suppose that the central bank has a model for this vector such that

$$E[Y_{t+h}|Y_{t-}] = f_h(Y_{t-}) \qquad h = 1, \ldots, H,$$

where Y_{t-} is the history of Y up through $t-1$ and E is the expectation of the central bank, and the f's are functions satisfying no other restrictions than those implied by the mean of expectation. We call this the *unconditional* forecast of Y_{t+h}, and write,

$$Y^u_{t+h,t-} = f_h(Y_{t-})$$

Sometimes we will drop the $t-$ subscript; whenever this second subscript is omitted it should be taken to be $t-$.

Unconditional refers to the fact that the forecast is conditioned only on actual information of the central bank at the time of the forecast. In particular, it is not conditioned on any counterfactual past or future values of any of the variables.

One variable in Y_t is the policy interest rate, i_t. We assume that the central bank has a model that gives forecasts conditional on certain presumed paths of the policy interest rate. Defining the conditioning path, $i^c_{t,H} = (i_t, i_{t+1}, \ldots, i_{t+H})$, we write the conditional forecast as,

$$Y^c_{t+h} = f^c_h(Y_{t-}|i^c_{t,H}) \qquad h = 0, 1, \ldots, H.$$

Our stylized view of forecasts in inflation reports is that they report conditional forecast paths for, say, two variables in Y_t (e.g., output growth and inflation) under one or more sets of conditioning assumptions, $i_{t,H}$. Our focus is on what agents can learn about the model (f_h) and the unconditional forecast Y^u_{t+h} from these conditional forecast paths.

Perhaps the most important point of the paper is that it takes a great deal of additional structure beyond what we have placed on the problem so far for the public to make any reliable inferences about f_h or Y^u based on the conditional forecast paths. To make the point, we lay out what we consider to be a "best reasonable case" scenario for the usefulness of conditional forecasts.

First, assume that the unconditional forecast is generated by simply iterating the one-step forecast:

$$Y^u_{t+h,t-} = f_1(Y^U_{h-1,t-}).$$

where $Y^U_{s,t-} = (Y^u_{t+s}, \ldots, Y^u_t, Y_{t-})$—this is the actual history of Y up through $t-1$ with the forecast up to $t+s$ appended. We now dispense with the 1 subscript on f.

Now we assume a simple structure for the conditional forecast. Intuitively, we simply drop the i equation from the model, replacing it with the exogenous path. Notationally, this requires that we partition the vector function f, separating the policy rate equation:

$$i^u_{t+h} = f_i(Y^U_{h-1,t-}) \tag{11}$$

$$Z^u_{t+h} = f_z(Y^U_{h-1,t-}) \tag{12}$$

where Z is all of Y except i, and now the subscript on f indicates which bit of f we are referring to. The conditional, or "replace the i equation," forecast is given by

$$i^c_{t+h} = \bar{\imath}_{t+h} \tag{13}$$

$$Z^c_{t+h} = f_z(Y^C_{h-1,t-}) \tag{14}$$

where $Y^c_{t+h} = (i^c_{t+h}, Z^c_{t+h})$, and $Y^C_{s,t-}$ is defined analogously to Y^U.

So long as f_z is defined for the resulting paths of Y^C, no complications arise in forming this thing we are calling a conditional forecast. It is in attempting to interpret and draw inferences from Z^c that the complications arise. Note that we defined Y^u to be an expectation (conditional on Y_{t-}) but we do not define Y^c to be any sort of conditional expectation. For the most part, it will not be important to do so.

We can now set about deriving the information content that Y^c has for Y^u. To do so, it is useful to consider a forecast generated by a one-period deviation from equilibrium play in the policy rate followed by return to equilibrium thereafter. That is, we have a unit shock, $i^c_t = i^u_t + 1$, and allow i to follow the i equation thereafter. As we have set things up, $Z^c_t = Z^u_t$: there is no contemporaneous reaction. This feature is inessential to the results.

$$Y^c_{t+h} = f(Y^C_{h-1,t-}) \quad h = 1, 2, \ldots, H.$$

Without loss of generality, we can define the difference between the conditional and unconditional forecasts to be,

$$g_h(Y_{t-}) = i^c_{t+h} - i^u_{t+h} \tag{15}$$

$$B_h(Y_{t-}) = Z^c_{t+h} - Z^u_{t+h} \tag{16}$$

for $h = 1, 2, \ldots, H$, where $g_0 = 1$ and all elements of B_0 are zero. The notation recognizes that—given the generality of the model, so far—these effects could differ depending on the state of the economy up through $t-1$.

Finally, we assume that g_h and B_h are constants independent of Y_{t-}. Notice that if the model, f, is linear, then g_h gives the conventional impulse response of the policy rate to a policy shock and B_h gives the impulse response of the other variables in the system to the policy shock.

Now we are in a position to relate Y^c to Y^u for arbitrary conditioning paths. Focus on a single Z variable, z, with impulse response at horizon h given by b_h. Straightforward substitution reveals

$$z^c_{t+h} - z^u_{t+h} = \sum_{i=0}^{h} b_i \varepsilon_{h-i}, \tag{17}$$

where ε is defined by

$$\varepsilon_{t+h} = i^c_{t+h} - f_i(Y^C_{h-1,t}) \tag{18}$$

and

$$i^c_{t+h} - i^u_{t+h} = \sum_{i=0}^{h} g_i \varepsilon_{t-i}. \tag{19}$$

Equation (18) defines ε_s to be the difference between i^c_s and what the model would predict it to be given what has come before. As is clear in the second form of the i equation, (19), the policy shock at $t+h$ is any part of the difference between i^c and i^u that is not accounted for the response of policy itself to the earlier policy shocks. Then (17) says that the deviation of the conditional forecast z from the unconditional forecast at horizon h is simply the sum effects of each of the shocks up to time h—shock ε_{h-i} has effect $b_i\varepsilon_{h-i}$ at time h.

We can stack these equations for $h = 0, \ldots, H$, and write the resulting system of equations as,

$$z^c - z^u = B^*\varepsilon \tag{20}$$

$$i^c - i^u = G\varepsilon \tag{21}$$

where $z^c = (Z^c_t, \ldots, Z^c_{t+h})'$ and Z^u, i^c and i^u, and ε are analogously defined as the $h+1$ values stacked. Finally, $B^* = Q(b)$, b is the b_h's stacked and $Q(x)$ is the lower triangular Toeplitz matrix with x_i on the i^{th} subdiagonal; finally, $G = Q(g)$.

These equations imply,

$$z^c - z^u = B^* G^{-1}(i^c - i^u) \tag{22}$$

$$z^c - z^u = C(i^c - i^u) \tag{23}$$

where $C = BG^{-1}$. The second equation is the dynamic equivalent of (9).

One more step is useful in interpreting this equation. Since B and C are lower triangular and Toeplitz, so is BC^{-1}. In particular, the final equation can be written,

$$z^c_{t+h} - z^u_{t+h} = \sum_{j=0}^{h} c_j (i^c_{t+h-j} - i^u_{t+h-j}) \tag{24}$$

where the c's are defined by, $c_0 = g_0$, and

$$c_h = g_h - \Sigma^h_{j=1} c_{h-j} g_j$$

As in the static case, at horizon zero the difference between the conditional and unconditional forecast is given by the conventional impulse response to a money shock times the deviation of the interest rate path from the unconditional.

Comparing (24) to (17) shows an important difference at larger horizons. The difference between z^c and z^u can alternatively be expressed as a sum of the conventional impulse response to a policy shock time the implied policy shock (the ε's) or the sum of the c's times the difference between the i^c_t and i^u_t, where the c's are a convolution of the impulse responses of z and i to the policy shock.

These equations give two alternative ways to adjust an unconditional forecast to obtain a conditional one. They emphasize that, to be consistent with the model, one cannot simply treat $i^c - i^u$ as the policy shock and adjust the unconditional forecast using the impulse response to a policy shock. This approach neglects the historical response of policy itself to the policy shock, which will be implicit (or explicit) in the structure of the model. Rather, starting with $(i^c - i^u)$ one must adjust using the convoluted c coefficients or, alternatively, one must derive the policy shocks implicit in the path of $(i^c - i^u)$.

This reveals one of the many complexities of conditional forecasting. It is very difficult to find exactly how the conditional forecasts in inflation reports are created. The BOE has, perhaps, been most explicit in this regard (e.g., in the August 2004 *Inflation Report*), but the text provided does not resolve how the issue just discussed has been handled.[7] Of course, without knowledge of such issues, the public cannot hope to invert the mapping from the conditional to the unconditional forecast.

3.2.1. *Inverting the mapping from conditional to unconditional.* When C is known, the structure of the problem is simple: if the public knows i^u, then together with the announcement it can deduce z^u. Alternatively, if the public knows z^u it can deduce i^u. Of course, these statements require very nearly complete information on the part

[7] Similarly, the issue is not resolved in Goodhart's (2005) account and attempt to approximate the BOE procedure.

of the public. The public must know complex convolutions of the parameters of the model, C, and it must know a substantial part of the central bank's unconditional forecast, either z^u or i^u.

Of course, the case we are interested in is when the central bank needs to communicate about the most likely future course of policy (i^u). Further, since i^u feeds back on the economy, it is difficult to imagine a case in which the public knows z^u and C, but not i^u. Thus, the cases in which the mapping is invertible require nearly complete information on the part of the public and probably have no realistic analog.

As in the static case, if the central bank gives two or more conditional forecasts, C is implied:

$$z^{c_2} - z^{c_1} = C(i^{c_2} - i^{c_1})$$

Given the structure of C, $C = Q(c)$, we can equivalently write,

$$Z^{c_2} - Z^{c_1} = Q(i^{c_2} - i^{c_1})c,$$

implying $c = Q(i^{c_2} - i^{c_1})^{-1}(Z^{c_2} - Z^{c_1})$. Invertibility in this case simply requires that the two conditioning assumptions for the interest rate differ sufficiently.

Alternatively, as we discussed above, c can be estimated from a time series of forecasts. Thus, if the public understands enough about the structure of the problem, it can infer c either exactly or imprecisely. This conclusion is rather modest. The public can infer a set of parameters that the central bank could alternatively just report. Even if they know c, they cannot invert the mapping from conditional to unconditional without knowing most of the unconditional forecast already.

We now turn to analysis of the empirical properties of the forecasts actually published.

4. Empirical Analysis of Published Central Bank Forecasts

If the public is to benefit from the forecasts in inflation reports, then these forecasts must have a great deal of systematic structure. Otherwise the public is pretty much in the position of the bus passengers described in the introduction. While they want to know when the bus will arrive, they may not know what to make of the report that "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."

We examine whether these forecasts in practice seem to have sufficient regularity and efficiency properties to allow the public to infer important information. We find that the forecasts do not appear to have strong properties in this regard.

This may seem like a strong criticism, but we do not intend it to be. We emphasize we are putting additional structure on many results that exist in the literature, many of them appearing in work by the central banks or in work they have commissioned. We will focus on the forecasts by the BOE and Riksbank, with analysis of the Fed's Greenbook forecast included for reference. This latter forecast is not intended as a communication device with the public and is not released until 5 years after the date of the forecast. Each of these banks performs regular analysis of the forecasts and the inflation targeters commission external reviews of their work, including the forecasts (e.g., Leeper, 2003; Pagan 2003). Other purely outside research has been done, the most familiar probably being Romer and Romer's (2000) work showing that the Greenbook forecast is pretty good by conventional benchmarks and Sims's (2002) more recent work reaching a similar conclusion, while being more critical about a number of particulars.[8] 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 on 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 (e.g., by market rates or by serial correlation in the policy changes themselves). This hints at the sort of problems we find below.

4.1. **Econometric issues.** Faust and Wright (2005) start with the key relationship derived above:

$$Z^c - Z^u = C(i^c - i^u) \tag{25}$$

The key result in their paper is that, under the assumption that the central bank's *unconditional* forecast is efficient, a time series of these conditional forecasts—say, from successive inflation reports—can be used to estimate $C = Q(c)$. Further, the structure of the problem imposes overidentifying restrictions, so that the assumption of efficiency of the underlying unconditional forecasts is testable.

Here is the reasoning. Efficiency of the unobserved unconditional forecast, $z^u_{t+h,t-1}$, implies that the *ex post* forecast error is unpredictable given information in the central bank's or public's information set at $t-1$:

$$z_{t+h} = z^u_{t+h} + \nu_{t+h}$$

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

where the ν is orthogonal to the information available to the central bank at time $t-1$. Thus, if we substitute the *ex post* realizations in (25), and expand the matrix notation, we get

$$Z^c_{t+h} - Z_{t+h} = \sum_{j=0}^{h} c_h(i^c_{t+h} - i_{t+h}) + \nu_{t+h} \quad h = 0, 1, \ldots, H. \tag{26}$$

This is a set of equations in the *ex post* errors in the conditional forecast, which are all observed. Thus, we can treat this expression as a regression equation and estimate the c's. Of course, ν is correlated with the regressors. Under the assumption of efficiency of the underlying unconditional forecast, ν is uncorrelated with anything in the central bank information set at t. Thus, we can consistently estimate this equation using standard instrumental variables approaches.

In the current context, the familiar problem of finding good instruments may be somewhat eased by the structure of the problem. We need something known at time t that predicts the difference between the assumed policy path and the subsequent outcome. Of course, the standard path assumed is intentionally counterfactual. For example, a flat path is often used, even when a constant policy rate is not viewed as a likely outcome. In this case, policy expectations derived from the term structure of interest rates (or associated derivative markets) are likely to provide good instruments.

The problem is a bit more subtle than this, however. If we simply estimate the system equation-by-equation, we would need separate instruments for the H variables, $(i^c_{t+h} - i_{t+h})$, $h = 0, \ldots, H$. In quarterly data, we might be interested in, say, $H \geq 4$, and obtaining the required instruments could be a problem.[9] Imposing the cross-equation restrictions implied by the fact that the c's are the same in all the equations greatly mitigates this problem, but weak instrument issues may remain. Faust and Wright (2005) explore these issues and suggest one approach for taking account of them.

This estimation of c is importantly premised on the fact that the underlying unconditional forecast is efficient. There is a simple way to test this assumption. We can augment the equations in (26) with any variables, w_t, known at time t:

$$Z^c_{t+h} - Z_{t+h} = \sum_{j=0}^{h} c_h(i^c_{t+h} - i_{t+h}) + \gamma' w_t + \nu_{t+h} \quad h = 0, 1, \ldots, H. \tag{27}$$

[9] The difference between the policy rate assumption and the unconditional expectation may be about the same, especially after the first couple of horizons. Similarly, the term structure data may not have greatly different information about adjacent horizons.

The hypothesis of efficiency implies $\gamma = 0$. Note that this equation can be estimated by OLS: correlation of the i variables and ν will mean that the c's from this equation are not consistently estimated, but the γ estimates remain consistent and the test will have the standard properties.[10]

This result is of considerable independent interest. Many papers have examined the properties of the Greenbook forecast and the forecasts in inflation reports. It is typical to simply analyze the forecasts as if they were *unconditional* forecasts. The papers then discuss possible implications flowing from the fact that the forecasts are in fact conditional. Of course, if the forecasts are found to be deficient, it may simply be that the problem stems from the fact that the conditioning assumptions are deliberately counterfactual. In the case of the Greenbook forecast, papers have documented that it performs very well as an unconditional forecast. It is never quite clear that this is a good thing: should we expect a forecast based on a counterfactual policy assumption to perform well?

The test proposed in Faust and Wright (2005) cuts through these issues. It provides a valid test of efficiency of an unconditional forecast based on the data for associated conditional forecasts. We emphasize that all these econometric results depend on some version of the assumptions given for the general model above. These impose a great deal of linearity and constancy through time in the assumed forecasting structure.

These assumptions are certainly not exactly true in practice. Whether they are even approximately true may be open to question. Problems with these assumptions only strengthen the main point of this paper, however. Our main point is that it is very difficult to learn about the unconditional forecast for the policy rate and the economy from conditional forecasts. If the model used by the central bank does not satisfy the sort of constancy and linearity assumptions we have made above, then the problem of inverting the mapping from the unconditional to the condition forecasts is surely made more difficult and quickly becomes entirely intractable.

4.2. **The 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.[11] 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,

[10] It is an empirical question whether the test based on OLS or IV estimates will have greater power.

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

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.[12]

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.[13]

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 in computing forecast errors. People differ on the appropriate approach.[14] 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.

5. Efficiency tests

For all three sets of forecasts, we compute two efficiency tests based on OLS regression of (27) with a single w variable added. In one case, denoted model f, the variable w is the forecast itself. In the other case, denote model i, the variable is the difference between the policy interest rate implied by futures markets three months hence and the current policy rate. A constant is also included and its value should be zero under forecast efficiency.

[12]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 used in constructing the Greenbook forecast.

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

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

Tables 1a, b, and c show the results for the BOE, Riksbank, and Greenbook, respectively. Consider the BOE results. For model f of inflation, at the longer horizons, we obtain border line rejections of efficiency near the 10 percent level. The joint test that α and γ are both zero is rejected at between the 5 and 10 percent levels. We do not see this as strong rejections of forecast efficiency.

The row labelled m-R^2 is the marginal contribution to the R^2 from including the w variable. This is a crude measure of the economic significance of the w variable: how much might incorporating this information help reduce the forecast error variance? These values are modest in the BOE inflation case.

The panel on GDP confirms the BOE result that the GDP forecast is not great. Efficiency is strongly rejected and the marginal R^2's are quite large.

The results for the Riksbank (Table 1b) are similar. There is somewhat stronger evidence against efficiency of the inflation forecast and somewhat larger marginal R^2's. There is quite strong evidence of inefficiency of the output forecast.

For the Greenbook (Table 1c), our results are incomplete and will be included in a later draft.

Overall, we see these results as modest-to-strong evidence against the efficiency of the central bank forecasts. Inefficiency severely complicates the public's inference problem.

Our theoretical result relates $z^c - z^u$ under the assumption of efficiency. It tells the public how to update its unconditional forecast when it observes the central bank's conditional forecast. Of course, if part of this difference may be due to predictable error in the underlying unconditional central bank forecast, this inference becomes muddled. As we will see below, attempts to estimate C from these forecasts give muddled results, and this may in practice be due to inefficiency of the underlying unconditional forecasts.

Of course, inefficiency in the forecast may have implications for policymaking itself and not just for the communication of policy. It would be wrong to jump to conclusions in this regard, however. At the policy meeting, a full range of conditional and unconditional forecasts may be discussed. The policy makers have available to them the full information set of the central bank. Thus, one cannot draw any clear implications from our results about the conduct of policy itself.

Our results are intended to show that the public, which does not have available to it the full information set of the central bank, will find it difficult to learn about the central bank's information set from these forecasts if they embed inefficiency.

5.1. **BOE Estimates of C Implied by Pairs of Conditional Forecasts.** As noted in the discussion of (26), for the BOE we can calculate the C implied by the pair of forecasts published in each inflation report. Remember that the C's are a convolution of the conventional impulse response of policy and the variable in question to a policy shock. To distinguish these from conventional impulse responses, we call these M-responses, where the M stands for multi, or mixed, or mongrel.

One might wonder what these 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 M responses are shown in Table 2. The policy shock is normalized to have a unit effect on rates initially which rises for a period and then decays. The conventional response of output is zero on impact by assumption and then monotonically increase in magnitude over the 5 periods shown. The implied M response is not monotonic and largely flattens out after the first 3 periods. While the effect of the shock at lag 5 is larger than at all earlier horizons, this gets distributed among all the implied deviations of i from the baseline.

The M-responses implied by pairs of BOE forecasts is shown in Table 3. The main impression one gets is that the implied responses are quite variable. They change sign and magnitude a great deal. At the very least, this is evidence that the pair of conditional forecasts created by the BOE was not created in the manner implied by our assumptions.

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.

5.2. **Estimating C's from a Time Series of Forecasts.** Table 4 reports our GMM estimates of M-responses based on the forecasts of the three central banks. In the case of the BOE, we make estimates based on the constant rate (bar) assumption only. Of course, these estimates are premised on efficiency of the underlying central bank forecast and constancy of the implied C's. We have collected considerable evidence against these assumptions already, and should expect very mixed results.

We estimate the systems using a standard 2-step GMM approach. For instruments, we follow Faust and Wright (2005) in using the difference between the current policy rate at t and the forward rate for h months hence. Our "full" instrument set for the system using horizons up to h is a constant and the forward spreads for horizons zero through h. Our "small" instrument set includes only a constant and the forward spread for horizon 3. Details of the method are provided in Faust and Wright (2005).

We estimate separate systems for horizons 0 through h, $h = 0, 1, 2, 3, 4$. The pairs of r, p rows in the table are for estimation of separate systems.

For example, if all the forecasts share the sort of properties shown in Table 3 for the BOE, we would expect widely varying results and perhaps large standard errors. This is, indeed, what we find. Going through the results in greater detail is not of much practical importance. The lesson here is that while on could in principle learn about the central bank's view of the effect of policy shocks from the conditional forecasts, in practice not much can be discerned.

5.3. **A Practical Perspective on the Results.** Advocates of the view that producing a conditional forecast in inflation reports is informative have more-or-less assumed that useful information can be gleaned from them. Upon deeper reflection (or after slogging through lots of algebra) it should be clear that conditional forecasts are very complicated to analyze and that judgmental conditional forecasts are perhaps unlikely to have sufficient structure to allow clear inferences to be drawn.

Below we venture a few suggestions about how more structure could be applied to the process without sacrificing the benefits of judgment embedded in the current system.

At the heart of all the reasoning about forecasts, however, is the view that we can form forecasts of output and inflation that are at fairly efficient. This view requires a bit of further examination.

As noted above, several papers have found that central bank forecasts perform well as viewed against standard benchmarks such as simple AR(1) forecasts or published forecasts from the private sector.

Sims (2002) conjectures that this result is mainly due to the resources the central banks commit to understanding the current state of the economy at the time of the forecast. That is, sector experts glean a great deal about what is happening, and this informs the jumping off point of the forecast and, thus, could improve the forecast over many horizons (until the initial conditions become irrelevant). If this is the main source of the quality of central bank forecasts, it suggests that structural models of the evolution of the economy going forward are not contributing much to the quality of the forecast.

To provide some tentative evidence on this view we consider the following simple-yet-sophisticated benchmark (SYSB) forecast. We begin with the central bank forecast of growth or inflation in the period of the forecast (the $h = 0$ forecast). This should embed all the central bank's special knowledge about the current state of the economy. Further, since most people think policy acts with a lag, the unconditional and conditional $h = 0$ forecasts should be very close. Our forecast path for

$h = 1, \ldots, H$ is then given by

$$(x_{t+h} - \mu_t) = \rho(x_{t+0} - \mu_t).$$

The forecast simply decays at the rate ρ to a mean chosen in real time based on the data available at t.

This is a two parameter forecasting model we will denote SYSB(ρ,μ), where μ is an indication of how the mean is chosen. Our goal is not to choose the best ρ and μ and show that we can beat the central bank forecast. Rather, our goal is to see if the properties of this simple forecast are relatively insensitive to the choice of μ and ρ. We would take this as evidence in favor of Sims's conjecture that the good properties of these forecasts stem mainly from study of the initial conditions.

We show results for all choices of ρ between zero and one. For μ we consider a simple sample mean of the last available observations at the time of the forecast. We call this parameter "window" and choose values from 1 to 40 (1 quarter to 10 years of data).

For inflation, this approach to the mean has an obvious problem. As of the beginning of the forecast samples, 1998 for the BOE, inflation was still adjusting to BOE independence. Thus, a mean of 10 years of past data probably was not a good estimate of inflation expectations going forward. Thus, we also consider setting μ_t equal to the last available observation on 6-10 year inflation expectations taken from consensus surveys. The longest central bank forecast we model is 6 quarters, so we view our measure from consensus forecasts as a reasonable estimate of some long-run mean of inflation that should not be changing much due to short-term considerations captured by the central bank.

Figure 1a reports the inflation results for the BOE-mkt forecast. The first column gives the mean absolute forecast error of the SYSB forecast relative to the BOE at 3 horizons (2, 4, and 6 quarters). Green indicates that the SYSB error is smaller; for red it is larger. More intense color indicates a larger relative gap.

In general, the BOE forecast outperforms the SYSB forecast. By the 6 quarter horizon, though, there is a substantial range of ρ and window size for which the SYSB forecast is competitive. Typically in forecasting the main concern is in trading off bias versus variance. The second two columns give the relative performance on bias and standard error, respectively. At the longest horizon, the SYSB may outperform the BOE on bias or standard error, generally not both. This result more or less confirms the efficiency results that this is a pretty good inflation forecast.

Figure 1b gives the SYSB forecast based on ρ and the μ taken from consensus forecasts. Since this gives only one measure of μ we only have one parameter that varies and the figure is easier to understand. Except at horizon 2, the SYSB forecast outperforms the BOE in absolute error, bias and standard error for a large range of smallish ρ's. (Note: two of the bias lines are actually off the bottom of the chart. This will be clarified in future drafts.) The interpretation may be that if one gets the long-run mean more or less correct, a simple forecast of AR(1) regression to the mean performs quite well relative to more complicated approaches.

Figure 1c gives the results for output. In this case, there are a substantial range of values of the window and ρ for which the SYSB forecast outperforms the BOE. The BOE forecast does better relative to the SYSB in this case as the forecast horizon grows.

The results for the Riksbank are very similar. The Riksbank forecast of inflation looks good relative to the SYSB(ρ, window) forecast, but the SYSB forecast based on the consensus μ outperforms the Riksbank. The Riksbank GDP forecast performs less well than the BOE relative to the SYSB.

Finally, the results for Greenbook are still in process. Tentative results are much like those for the other two banks.

Overall, these results tend to support Sims's conjecture. Of course, like Sims and the earlier papers, these results treat the central bank forecasts as if they were unconditional. Future versions of the paper will deal with this problem more explicitly. A more complete analysis of the SYSB forecasts is in order. Thus, these results are just suggestive.

Still these results add a cautionary note to our practical suggestions. Forecasts that ignore essentially all information except the initial condition as assessed by the central bank seem to perform very well. This suggests that our ability to analyze how different factors affect economic variables over a forecast horizon may be modest. From the standpoint of this paper's emphasis on conditional forecasts as communication devices, this provides a stronger case that conditional forecasts may be difficult to understand. Of course, there may also be implications for the policymaking process, but these are not the focus here.

6. Practical Implications

This section draws out the practical implications of the analyses in the text.

6.1. **Making Forecasts.**

While our goal is mainly to discuss how to improve the use of forecasts in communication, one important step in this process would be improving the properties of the forecasts themselves. This might have implications for policymaking as well as communication.

We do not plan to say much about the implications for policymaking. Nothing in our analysis precludes having analysts present a menu of conditional policy paths and their likely impacts to their policy committees. Policymakers select a path from the menu or iterate with the analysts to find the most desirable path. Under the assumption that policy is at least in some limited sense optimizing, *by definition* the selected policy path and the implied paths for macro variables become the central bank's unconditional forecast. That forecast reflects all information available to policymakers, including any signals they received and judgments they may have formed about the state of the economy, the current state of their preferences and technology of policy selection, and so forth.

One clear suggestion of our analysis is that (absent other considerations) maximizing transparency suggests that central banks would do well to start their communication with an unconditional forecast that includes the future path of the policy instrument. To our knowledge, no central bank now does this.[15] In addition to announcing what it regards as the most likely path for the economy and policy, an unconditional forecast, by embedding equilibrium policy behavior, can shine light on the central bank's reaction function. Because the unconditional forecast reflects any private information the central bank may possess, it is a mechanism for shedding crucial aspects of information, which enhances transparency. While this surely would not shed all information, in our view, it would share the most important aspects of private information that central banks now retain since the "transparency revolution."

It may be useful for the central bank to generate, if not publish, a benchmark statistical forecast. This would serve as a guide in the sense that both judgmental adjustments and conditional projections can be built up from it. judgmental adjustments are applied to a statistical forecast to bring the central bank's expert analysis to bear. We suspect that part of the problematic properties of conditional forecasts that we isolate is that it is difficult for sector experts to consistently condition their judgmental advice on an arbitrary path of the policy rate. Doing so requires a general equilibrium model with a level of detail that no current models begin to approach. In practice, we suspect that the judgmental sector input may be biased toward the unconditional forecast. A statistical forecast, then, would be just a first step in constructing the published unconditional projection and building up any conditional projections.

With a statistical benchmark forecast in hand, central bank analysts can gauge how likely judgmental adjustments are relative to historical forecast errors. To assess the reliability of forecasts, model users need to know how much judgment is moving the statistical model away from past patterns of correlation among variables. Large

[15]The Reserve Bank of New Zealand might be an exception, though it does not characterize its central projection as an unconditional forecast.

judgmental adjustments in the absence of large economic events can be a symptom of model breakdown.

Conditional projections can be built up from unconditional ones by solving for the sequence (or sequences) of shocks necessary to produce the conditioning set.[16] In fact, any method of constructing conditional projections requires perturbing the unconditional projection in certain ways. To gauge the reliability of the conditional forecasts, it is important to understand and examine the requisite intervention on the unconditional forecast.[17]

It is easier and less controversial to produce unconditional forecasts. Constructing economically meaningful conditional forecasts requires identifying assumptions. These assumptions are inevitably contentious, could even be wrong. Moreover, we suspect that communicating to the public clearly about the nature of the identifying restrictions underlying conditional projections may be a fool's errand. Central banks have not pursued this errand thus far, but without such information it is difficult to know what to make of conditional forecasts.

Finally, unconditional forecasts—both with and without judgmental adjustment—can be readily checked for consistency with data using conventional forecast performance tests. Conditional forecasts of the type most often published by central banks make forecast evaluation more difficult.

6.2. **Communication.** The most straightforward way for the central bank to communicate to the public is to report its estimate of the most likely paths the policy rate and policy's goal variables will follow. Because the central bank's judgment is an important source of private information, it is probably useful to communicate the judgment and how it affects the outlook. It is much easier to communicate the role of judgment in an unconditional forecast, which is not muddied by a counterfactual assumption about policy behavior. We have focussed on the expected path rather than higher moments, because with think that this is most important. Clearly, forecast uncertainty is also important, but this paper has nothing to add on this topic.

There may be sociopolitical reasons why the unconditional forecast and especially the forecast of policy rates should not be communicated. Thus, it is worth mentioning some suggestions that fall short of this more complete transparency.

It is wise for central banks to avoid grossly counterfactual conditioning assumptions. At times a constant interest rate assumption is grossly counterfactual, and such counterfactuals almost certainly put a strain on our ability to sensibly forecast. We probably should avoid attempting to comment on the implications of gross

[16]Waggoner and Zha (1999) describe the mechanics of doing this.

[17]Leeper and Zha (2003) discuss this in detail with regard to monetary policy interventions.

counterfactuals. Banks that routinely condition on market interest rates probably successfully avoid gross counterfactuals. Those that always condition on a constant interest rate over the forecast horizon probably do not.

Certain problems emerge in interpreting what the central bank thinks about the market-based path. The most natural thing would be for the central bank to state whether it finds this path likely or not, but we are ruling this out. The conundrum of transparency in this area is that the central bank wants the public to know where rates are going, but does not want to tell them. Moving to the market-based path may seem to put the central bank in the position of lending credibility to that path. Perceiving this, some central banks appear to distance themselves from the market path.[18]

It seems counterproductive for the central bank to publish a conditional forecast—using, for example, market futures interest rates—and then to distance itself from the conditioning assumption by asserting that it does not reflect the views of the policymakers. By disowning the conditioning path, the inflation report leaves the reader wondering why that particular path of interest rates is especially worthy of attention in the report, particularly when the path in question constitutes the central projection.

It is worth emphasizing that conditional forecasts specifying alternative scenarios for variables other than policy can be useful for showing the effects of monetary policy and for explaining policy choices. A recent *Monetary Policy Statement* by the Reserve Bank of New Zealand serves as a nice illustration. The RBNZ's central

[18]This pattern of publishing and distancing is followed by the Bank of England and the Riksbank. The BOE has been publishing forecasts conditional on both a constant rate and market rates since 1998, and recently has adopted the market-rate conditioning path as its central projection. Each inflation report contains a disclaimer similar in spirit to this one: "Even if market participants and the MPC have a common view about the economic outlook, the MPC may decide that it is appropriate for official interest rates to follow a different path from that implied by the yield curve in order to achieve its *Remit*." (Iinflation Report, August 2004, p. 43)

The Riksbank only recently began to condition on an alternative, market-based path for the repo rate. It gives the following explanation for the alternative assumption: "The aim to be as open and clear as possible makes it natural to publish this additional information. It should be emphasised that this does not mean that an interest rate trend that follows forward rates should be interpreted as the monetary policy line the Executive Board considers most probable." (Inflation Report, 2005:1, p. 5)

In contrast, in its 2005:1 inflation report, the Norges Bank writes that forward rates "provide a reasonable indication of [interest rate] expectations in the period to 2007." (p. 9) Presumably, "reasonable" means the market view coincides with the Norges Bank's view. Remarkably, over the period 2007-2008, the Norges Bank adopts a view of the interest rate path that differs from market views.

projection assumes a path for interest rates that reflects the Governor's views. Off of this path the bank constructs inflation and output projections. In March 2005, the *MPS* includes two additional alternative scenarios—an upside and a downside—with specific structural stories attached to each. The impacts of these scenarios are projected conditional on the interest rate path for the central projection. Naturally, in one case the inflation projection is significantly higher than in the central projection, while in the other it is lower. A footnote explains that if either scenario were to occur, policy would respond, leading to different outcomes for inflation and GDP. The counterfactual shows that the interest rate path is state dependent and it elaborates on specific contingencies under which the RBNZ would make different policy choices. Arbitrary conditioning paths—such as a constant interest rate or market interest rates—communicate almost nothing about the central bank's intentions and very little about the central bank's information.

6.3. **Objections.** It is not hard to imagine a number of objections that might be raised to having central banks announce the most likely path of future policy. Some objections are specific to this proposal; others are the ones that arise whenever any marginal increase in transparency is contemplated.

One objection is that the proposal requires too much information. After all, how can even a central banker be expected to know what the interest rate will be a year or two from now? The proposal simply calls for clear communication of the central bank's current understanding of likely outcomes. If the public already has the same view as the central bank about the path of policy, then the goal has been achieved. If, however, the central bank has a better view of the policy interest rate than the public, it should share that view with the public. An unconditional projection of the rate accomplishes this.

More generally, it is important to recognize that producing an unconditional forecast requires less, not more information than conditional ones. As we argued above, unconditional projections are easier to construct and, because they do not require taking as strong a stand on the structure of the economy, may actually take less information to form.

If inflation reports repeatedly report interest rate paths that ultimately are not realized—because shocks are continually hitting the economy and pushing rates off track—it will hinder communication and credibility. The central bank will be forced to spend resources explaining why it deviated from the unconditional path and why those deviations do not constitute a change from stated policy. What's wrong with that? Explaining interest rate decisions is a crucial ingredient to transparency. But we suspect these concerns are overstated. To our knowledge, central banks whose central projection is an unchanged interest rate are not routinely accused of reneging on previous promises. Deviations from an unconditional path are an inevitable consequence of uncertainty and communication about the nature of uncertainty can never be avoided.

Finally, we come to the big objection: reporting an interest rate path would for sociopolitical reasons compromise the central bank's ability to conduct the proper policy. This may be true. We are not experts in this area and will make only one comment. It is important that those trading off the benefits of more transparency versus the political costs be well aware of both sides of the equation. This paper mainly fills out the benefit to transparency.

7. Conclusions

Individuals in the private sector care about where the policy rate is headed because it directly affects their choices and because of its affect on the economy more generally. This paper examines the ways that forecast information can help the central bank communicate its views of where policy is headed.

It is natural to focus the discussion around forecasts: conditional forecasts have an essential role in the monetary policy formation process. In the process of determining the best policy choice at policy meetings is easiest to cast as a choice among various conditional outcomes posited to ensue if a given various policies are chosen. Conceptually, a wide range of these can be discussed and evaluated in light of all the information the policy board has at its disposal.

These facts do not imply that conditional forecasts as they are typically reported in inflation reports are of much value in communicating the central bank's views. Even in the best of cases, one or two conditional forecast paths for one or two variables conditioned on arbitrary interest rate paths that may be far from the optimum are in fact quite complicated to interpret.

The communication problem is magnified if the central bank attempts to distance itself from any claim that the paths conditioned on reflect what might actually happen to the policy rate.

If we consider only the goal of transparency, there are some simple steps the central bank can take. A great deal of information would be provided in an unconditional forecast of key variables and of the policy rate. This may not be politically acceptable. This paper provides some other suggestions about how the use of forecasts may be improved even absent the most direct approach.

References

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

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

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

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

[5] 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.

[6] 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.

[7] 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.

[8] 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.

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

[10] 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.

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

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

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

[14] 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.

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

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

[17] Norges Bank (2005), *Inflation Report*, March.

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

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

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

[21] 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.

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

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

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

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

[26] 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.

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

Table 1a. Efficiency tests of BOE forecast

	forecast horizon				
	0	1	2	3	4
			inflation		
			model: f		
α	0.42	0.81	0.78	0.90	0.52
p	0.12	0.02	0.03	0.02	0.18
β	-0.10	-0.17	-0.18	-0.19	-0.17
p	0.25	0.13	0.19	0.07	0.13
F-p	0.24	0.04	0.06	0.04	0.28
m-R^2	0.06	0.09	0.06	0.07	0.09
			model: i		
α	0.17	0.48	0.37	0.46	0.03
p	0.45	0.07	0.11	0.14	0.87
β	0.26	-0.26	1.05	-0.03	0.79
p	0.42	0.43	0.02	0.95	0.05
F-p	0.43	0.13	0.02	0.26	0.11
m-R^2	0.03	0.01	0.15	0.00	0.09
			output growth		
			model: f		
α	2.85	2.45	2.15	2.86	2.00
p	0.00	0.00	0.00	0.00	0.00
β	-1.08	-1.08	-0.91	-1.23	-0.78
p	0.00	0.00	0.00	0.00	0.00
F-p	0.00	0.00	0.00	0.00	0.00
m-R^2	0.65	0.48	0.29	0.46	0.23
			model: i		
α	0.49	0.47	0.36	0.65	0.29
p	0.20	0.23	0.37	0.06	0.44
β	0.46	-0.77	-0.53	-1.38	-0.16
p	0.44	0.07	0.31	0.01	0.73
F-p	0.27	0.07	0.41	0.00	0.72
m-R^2	0.03	0.06	0.03	0.14	0.00

Note: See the notes at end of Table 1.

Table 1b. Efficiency tests of Riksbank forecast

	forecast horizon				
	0	1	2	3	4
			inflation		
			model: f		
α	0.60	0.83	0.48	0.25	1.11
p	0.29	0.06	0.24	0.50	0.03
β	0.02	0.02	0.05	0.02	0.16
p	0.53	0.38	0.26	0.22	0.00
F-p	0.06	0.00	0.02	0.15	0.00
m-R^2	0.01	0.03	0.11	0.02	0.26
			model: i		
α	0.23	1.40	1.01	-1.38	0.75
p	0.71	0.17	0.28	0.04	0.08
β	1.00	-1.36	-1.13	4.16	0.23
p	0.32	0.44	0.60	0.00	0.83
F-p	0.27	0.00	0.28	0.00	0.00
m-R^2	0.04	0.06	0.02	0.29	0.00
			output growth		
			model: f		
α	1.65	1.51	2.06	2.24	3.05
p	0.01	0.10	0.01	0.08	0.00
β	-0.89	-1.00	-1.25	-1.27	-1.49
p	0.00	0.00	0.00	0.02	0.00
F-p	0.00	0.00	0.00	0.00	0.00
m-R^2	0.66	0.59	0.52	0.30	0.22
			model: i		
α	0.68	-0.36	0.14	-0.50	-1.46
p	0.49	0.66	0.57	0.40	0.00
β	-1.48	-1.86	-3.05	-1.69	-0.06
p	0.30	0.31	0.00	0.15	0.84
F-p	0.52	0.06	0.00	0.00	0.00
m-R^2	0.05	0.09	0.27	0.13	0.00

Note: See the notes at end of Table 1.

Notes: In all cases the tests are based on a regression of the form $e_{ht} = \alpha + \beta x_t + \sum_{i=0}^{h} \gamma_i e_{iht} + \text{error}$, where e_{ht} is the horizon h forecast error for one quarter output growth or inflation in quarter t and e_{iht} is the *ex post* difference between the policy rate and the conditioning assumption h periods earlier for period t. The columns are for separate regressions at each horizon. For model f, x_t forecast for t from $t-h$; for model i, the forward interest differential from time t as described in the text. Row labelled p are the p-value from a conventional two-tailed t-test that the coefficient in the row above is zero. The F-p is the p-value for a test that both α and β are zero. The $m - R^2$ gives the marginal increase in the R^2 statistic from including the x variable. All statistics are based on Newey-West standard errors with lag length equal to the horizon h. All p-values are based on conventional asymptotics. The sample sizes are 1998:1-2003:2, 2000:2-2003:2, and Greenbooks from 1988/10/26 through 1999/12/15, for the BOE, Riksbank, and Fed forecasts, respectively.

Table 2. Impulse response of interest rate and output and implied M-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
M-resp: 100 c_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 3a. M-Responses for inflation implied by pairs of BOE Forecasts

hor:	0	1	2	3	4	5
1998:1	0.00	0.00	-0.06	-0.45	-2.09	-4.38
1998:2	0.00	0.00	-0.31	-0.69	-0.14	0.22
1998:3	0.00	0.00	0.00	0.00	0.00	0.00
1998:4	-0.28	0.85	6.71	11.15	12.16	7.39
1999:1	0.04	3.01	3.44	2.03	2.88	0.63
1999:2	0.00	0.00	-0.65	-1.44	-2.13	-2.71
1999:3	0.39	0.67	-0.71	-1.67	0.22	3.12
1999:4	0.00	-0.53	0.15	1.84	3.99	1.69
2000:1	0.00	0.22	-0.96	-3.09	-7.84	-9.86
2001:1	0.00	0.00	0.12	-1.26	-3.14	-3.81
2001:2	0.00	0.00	0.00	0.00	0.13	-0.46
2001:3	0.00	0.00	0.00	0.00	0.00	-1.09
2001:4	0.00	0.00	0.00	0.00	-0.67	2.00
2002:2	0.00	0.00	-0.47	-1.65	-3.12	-5.87
2003:1	0.59	-0.40	0.20	0.20	-0.40	-0.59
2003:2	0.00	0.00	0.00	0.11	1.25	0.00
2003:4	0.00	0.00	0.12	-1.36	-2.42	-4.60
2004:1	0.00	0.00	0.00	0.24	-0.30	-2.02
2004:2	0.00	0.00	0.00	0.12	-0.74	-2.98
2004:3	0.00	0.00	0.00	-0.86	-3.43	-6.85
2004:4	0.00	0.00	0.00	0.00	1.12	1.77
ave.	0.04	0.18	0.36	0.15	-0.22	-1.35
min.	-0.28	-0.53	-0.96	-3.09	-7.84	-9.86
max.	0.59	3.01	6.71	11.15	12.16	7.39

Notes: See the notes at end of Table 3b.

Table 3a. M-Responses for output growth implied by pairs of BOE Forecasts

hor:	0	1	2	3	4	5
1998:1	0.58	3.25	14.46	17.58	12.67	6.26
1998:2	1.83	6.86	10.05	7.72	5.78	5.63
1998:3	0.33	0.11	0.88	0.65	0.63	-1.46
1998:4	-0.85	4.84	9.17	4.71	-15.79	-49.08
1999:1	-0.02	-1.55	-4.13	-10.31	-21.95	-36.99
1999:2	0.00	-0.66	-2.12	-2.94	-1.45	0.42
1999:3	0.00	0.00	1.18	2.81	3.56	2.55
1999:4	0.00	0.00	1.05	1.28	-1.47	-8.95
2000:1	0.00	0.22	-0.29	-6.14	-16.31	-29.05
2001:1	0.00	0.23	-2.27	-8.56	-16.27	-25.40
2001:2	0.00	0.13	-0.33	-1.77	-4.43	-6.37
2001:3	0.00	-1.10	-2.12	-5.32	-9.63	-16.61
2001:4	0.00	-0.67	0.67	-0.67	-2.11	-4.77
2002:2	0.00	-0.23	-2.11	-9.06	-24.71	-49.98
2003:1	-0.59	-0.19	0.19	0.20	-0.19	-0.58
2003:2	-0.11	-1.48	-2.97	-5.30	-6.82	-7.49
2003:4	0.00	0.12	-1.01	-6.04	-17.41	-35.22
2004:1	0.00	0.96	-2.13	-6.53	-15.28	-23.69
2004:2	0.12	-0.36	-5.81	-11.03	-17.63	-22.76
2004:3	0.00	-1.69	-5.87	-9.22	-13.44	-15.92
2004:4	0.00	1.11	1.76	1.76	1.76	3.97
ave.	0.06	0.47	0.39	-2.20	-7.64	-15.02
min.	-0.85	-1.69	-5.87	-11.03	-24.71	-49.98
max.	1.83	6.86	14.46	17.58	12.67	6.26

Notes: These are the M-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. A few IRs are missing; this occurs when the pair of interest rate paths does not differ sufficiently to meet the invertibility condition.

Table 4a. Estimated M-Response of inflation: BOE

h:	0	1	2	3	4	J-stat.
			inst set: small			
r	1.62	.	.	.	.	0.07
p	0.62	.	.	.	.	0.80
r	-0.83	-0.52	.	.	.	0.84
p	0.68	0.96	.	.	.	0.66
r	-2.13	10.38	-27.55	.	.	2.02
p	0.40	0.37	0.55	.	.	0.57
r	-1.65	5.98	-3.63	-19.11	.	2.59
p	0.40	0.29	0.48	0.26	.	0.63
r	-1.01	4.43	-5.18	0.34	6.79	3.17
p	0.50	0.29	0.25	0.90	0.70	0.67
			inst set: full			
r	-0.09	.	.	.	.	0.91
p	0.97	.	.	.	.	0.34
r	0.10	-3.43	.	.	.	1.53
p	0.93	0.66	.	.	.	0.67
r	-1.15	4.41	-3.34	.	.	3.74
p	0.13	0.08	0.74	.	.	0.71
r	-1.26	4.12	-1.98	-11.64	.	5.03
p	0.02	0.00	0.25	0.16	.	0.89
r	-1.60	4.64	-3.64	-0.64	3.64	5.16
p	0.00	0.00	0.00	0.02	0.40	0.99

Notes: See notes at end of Table 4.

Table 4b. Estimated M-Response of output: BOE

h:	0	1	2	3	4	J-stat.
			inst set: small			
r	-2.75	.	.	.	.	0.00
p	0.65	.	.	.	.	0.97
r	-4.14	6.44	.	.	.	0.09
p	0.16	0.75	.	.	.	0.96
r	-3.81	3.33	17.50	.	.	0.14
p	0.03	0.44	0.31	.	.	0.99
r	-3.34	5.35	-9.11	35.65	.	3.54
p	0.05	0.21	0.06	0.02	.	0.47
r	-3.58	6.14	-4.45	-2.99	25.94	4.03
p	0.04	0.16	0.18	0.44	0.07	0.54
			inst set: full			
r	3.08	.	.	.	.	0.61
p	0.56	.	.	.	.	0.43
r	-1.17	-9.34	.	.	.	3.68
p	0.14	0.36	.	.	.	0.30
r	-1.95	-0.45	22.09	.	.	4.04
p	0.01	0.86	0.20	.	.	0.67
r	-2.43	6.16	-13.67	43.29	.	5.08
p	0.00	0.00	0.00	0.01	.	0.89
r	-1.97	5.02	-6.75	-0.22	29.60	5.16
p	0.00	0.00	0.00	0.83	0.00	0.99

Notes: See notes at end of Table 4.

Table 4c. Estimated M-Response of inflation: Riksbank

h:	0	1	2	3	4	J-stat.
			inst set: small			
r	13.42	.	.	.	.	0.85
p	0.00	.	.	.	.	0.36
r	0.62	-5.22	.	.	.	2.66
p	0.90	0.83	.	.	.	0.26
r	5.71	-35.78	158.79	.	.	2.65
p	0.13	0.10	0.09	.	.	0.45
r	-2.32	4.70	-5.94	37.62	.	2.98
p	0.29	0.61	0.84	0.75	.	0.56
			inst set: full			
r	5.68	.	.	.	.	1.01
p	0.03	.	.	.	.	0.32
r	-14.50	66.61	.	.	.	2.58
p	0.15	0.20	.	.	.	0.46
r	-5.34	2.53	24.12	.	.	3.06
p	0.01	0.51	0.05	.	.	0.80
r	-6.27	5.10	16.92	-53.39	.	3.16
p	0.00	0.00	0.00	0.00	.	0.98

Notes: See notes at end of Table 4.

Table 4d. Estimated M-Response of output: Riksbank

h:	0	1	2	3	4	J-stat.
			inst set: small			
r	-15.40	.	.	.	.	0.01
p	0.21	.	.	.	.	0.91
r	4.13	-30.25	.	.	.	2.77
p	0.53	0.34	.	.	.	0.25
r	-12.25	62.95	-269.68	.	.	2.71
p	0.25	0.22	0.23	.	.	0.44
r	3.57	-17.47	59.48	-237.89	.	2.81
p	0.20	0.25	0.33	0.34	.	0.59
			inst set: full			
r	-10.06	.	.	.	.	0.04
p	0.19	.	.	.	.	0.85
r	1.31	-2.57	.	.	.	2.66
p	0.57	0.68	.	.	.	0.45
r	-2.17	30.77	-145.82	.	.	3.15
p	0.02	0.00	0.00	.	.	0.79
r	8.91	-13.95	11.56	-46.44	.	3.16
p	0.00	0.00	0.00	0.00	.	0.98

Notes: See notes at end of Table 4.

Notes: The tables report the results of estimating M-responses using a conventional 2-step GMM procedure. Details on the system are provided in the Appendix and Faust and Wright (2005). The full instrument and small instrument sets are described in the text. Each pair of rows, labelled (r,p) reports the results for a separate estimate of the system, where the system involves as many equations as there are horizons h reported in the row. The r row gives the estimated M-response at horizon h to the policy shock. The units are percentage points of M-response (on quarterly output growth or inflation) to a one percentage point shock. The row labelled p gives the associated asymptotic p-value for a t-test that the coefficient above is zero. The J-stat. column gives the conventional Hansen J test of overidentifying restrictions in the GMM context and the associated asymptotic p-value.

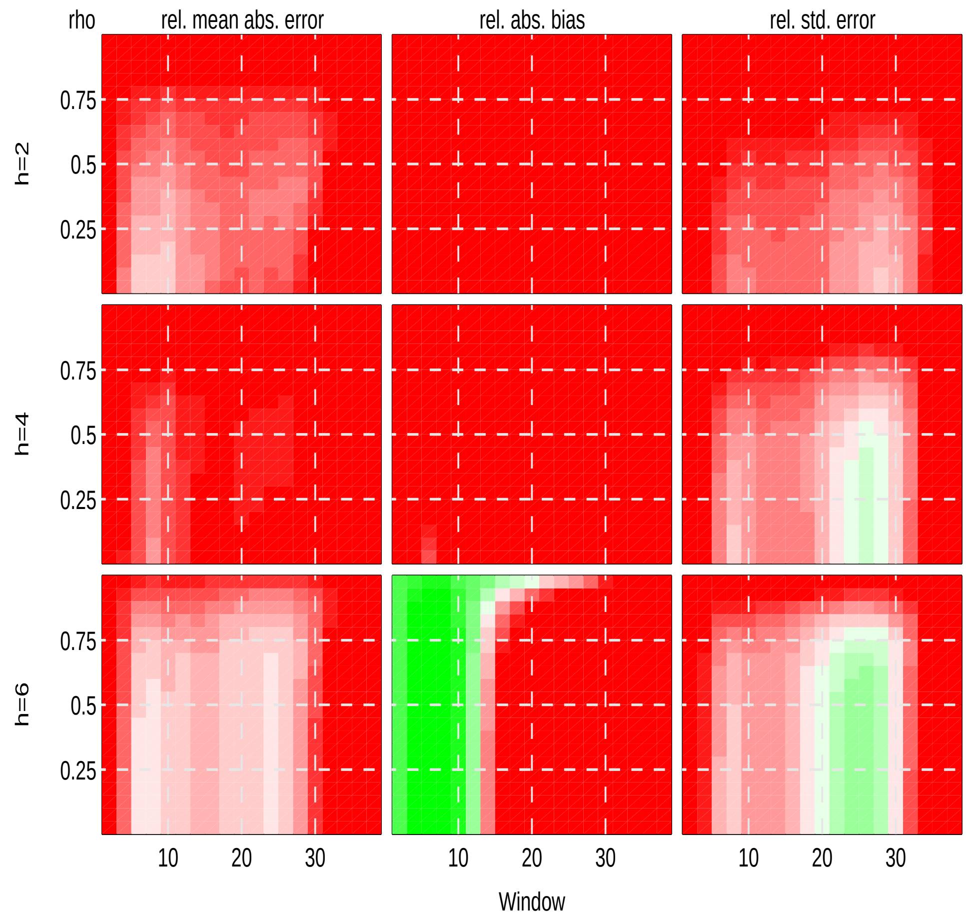

Fig. 1a: Performance of Inflation Forecast: SYSB(rho,window) relative to BOE−mkt

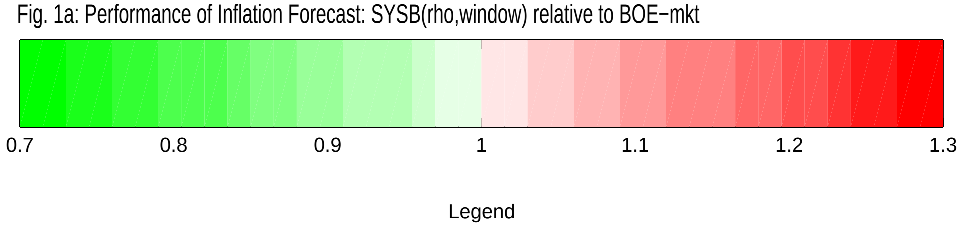

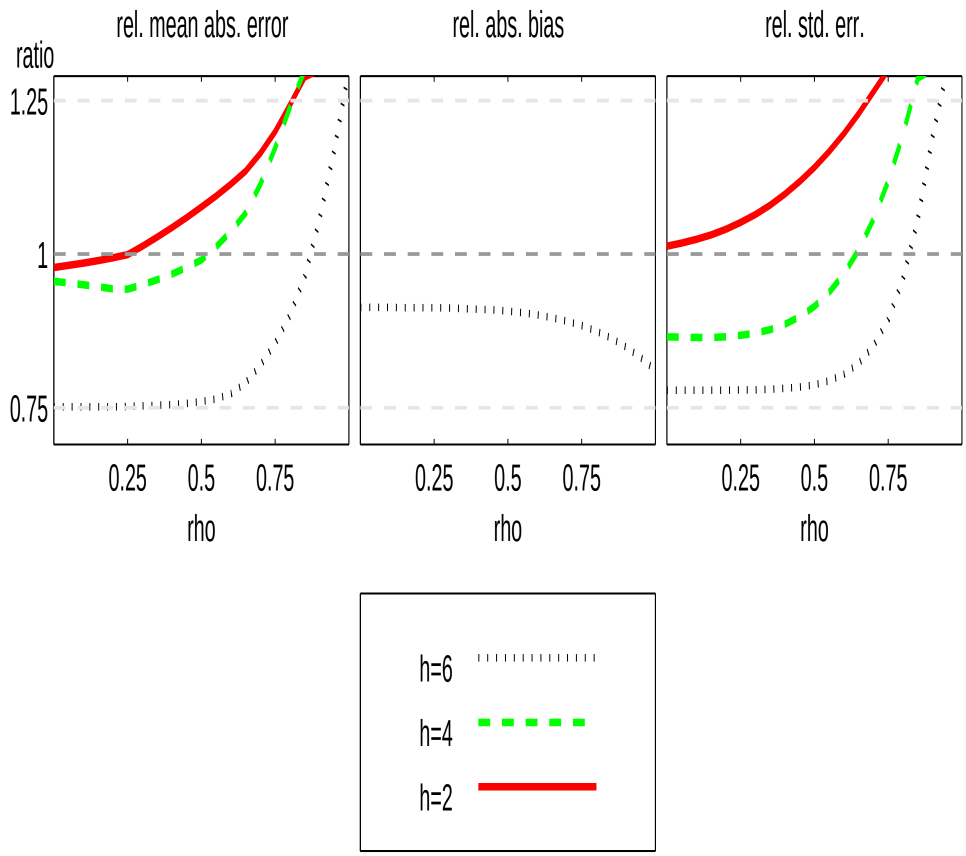

Fig. 1b. Performance of Inflation Forecast: SYSB(rho,consens) relative to BOE−mkt

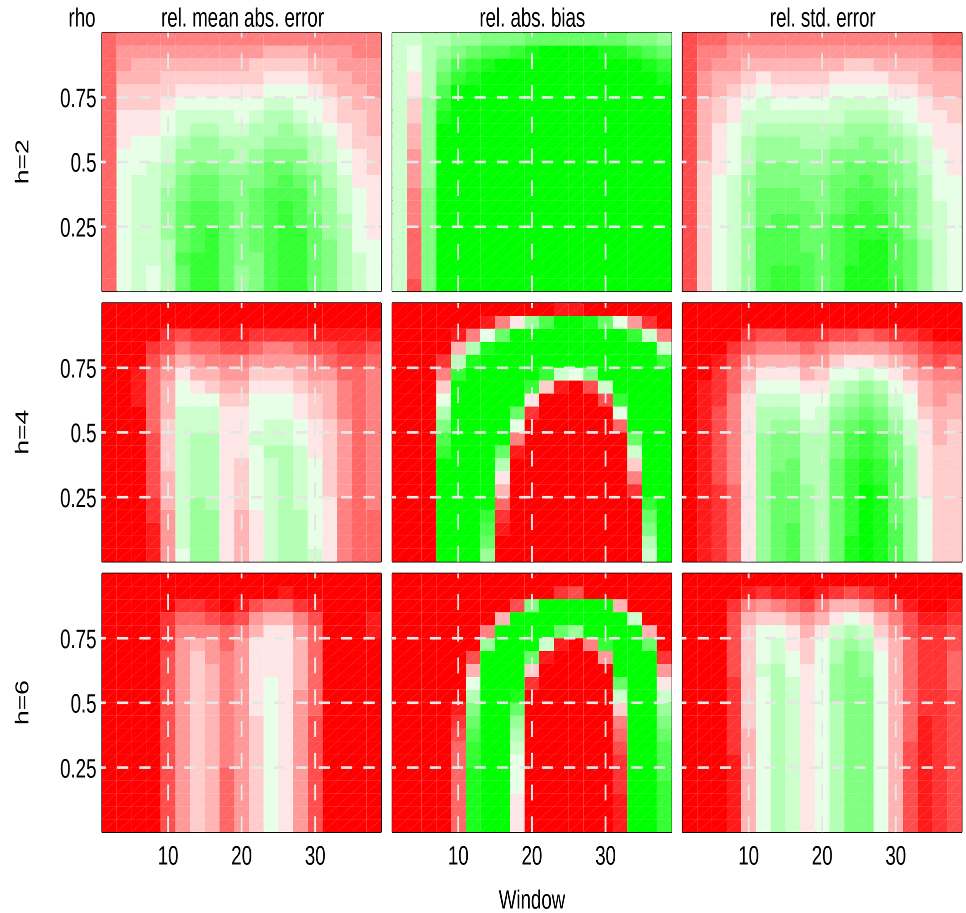

Fig. 1c: Performance of Output Forecast: SYSB(rho,window) relative to BOE−mkt

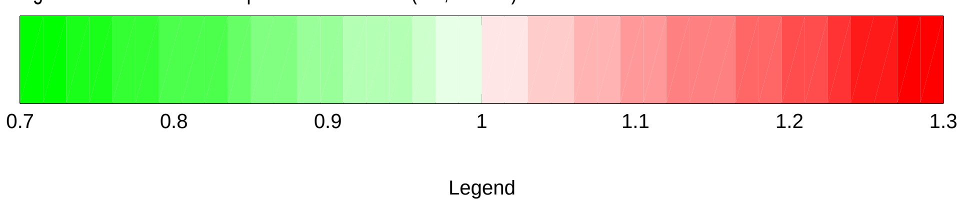

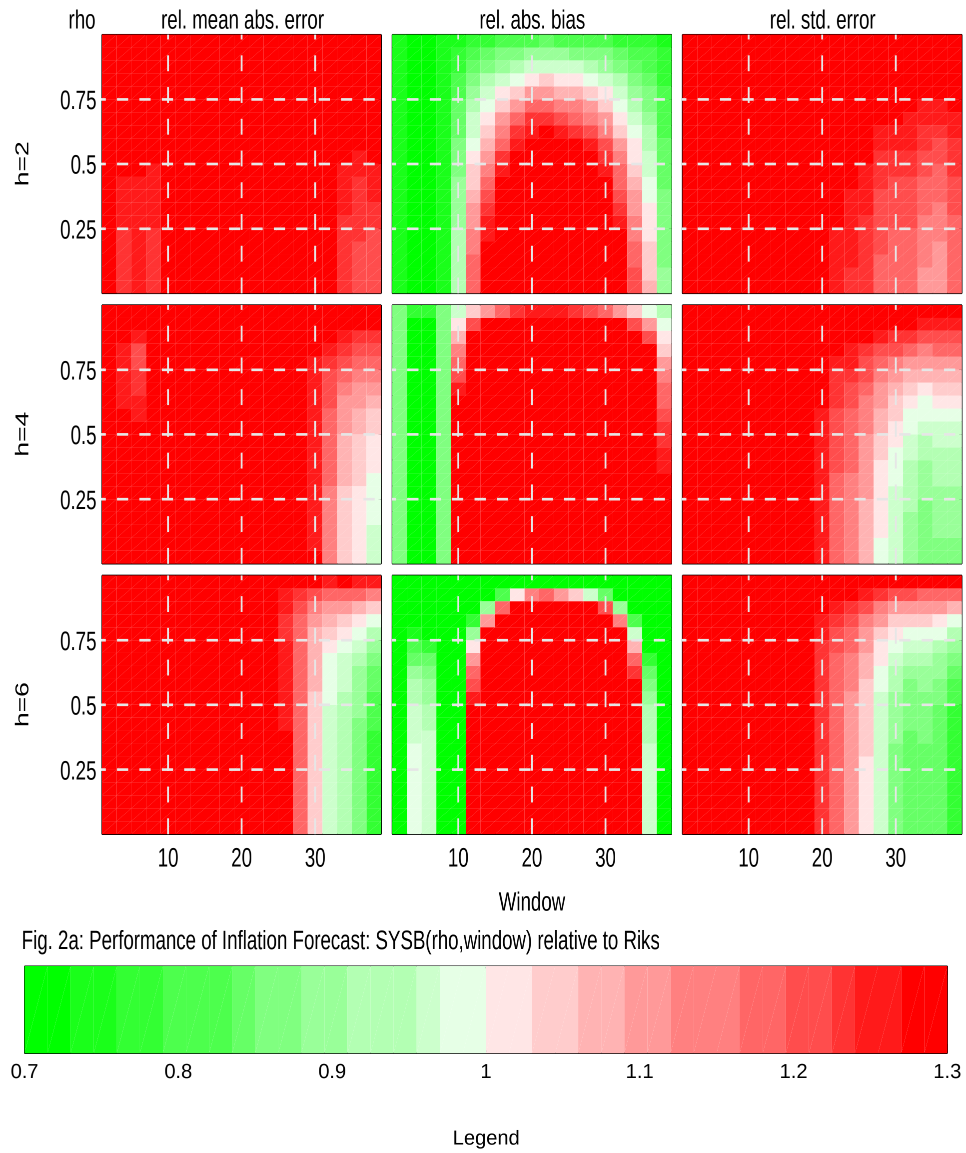

Fig. 2a: Performance of Inflation Forecast: SYSB(rho,window) relative to Riks

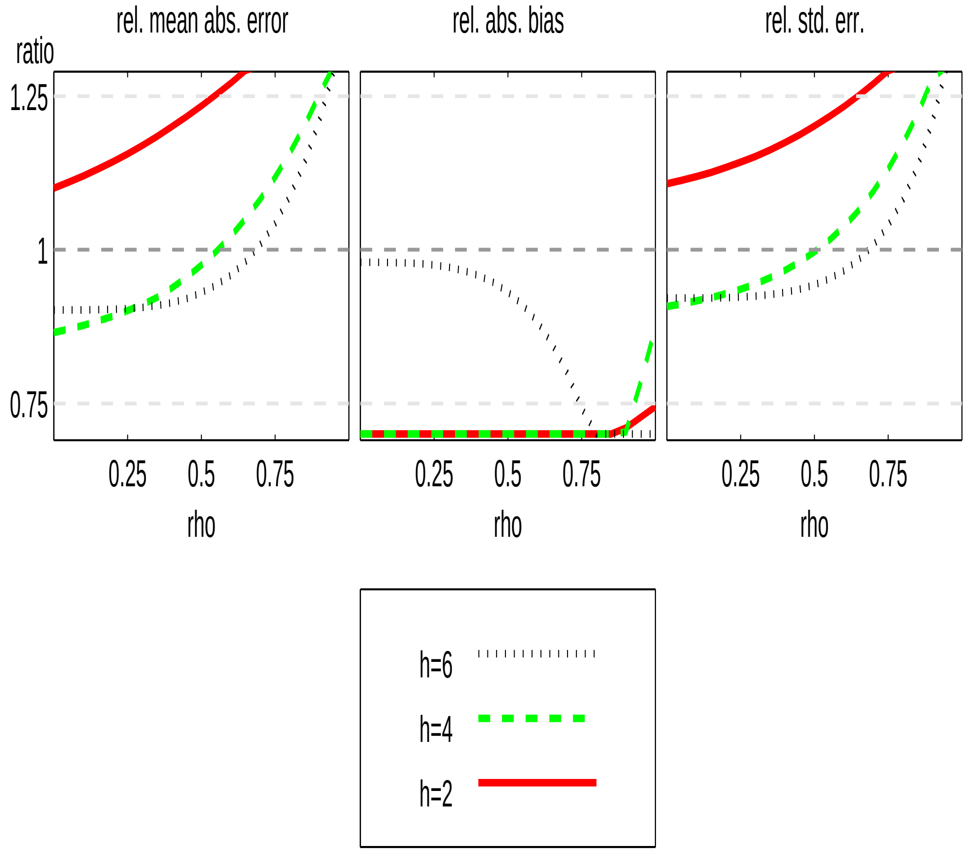

Fig. 2b: Performance of Inflation Forecast: SYSB(rho,consens) relative to Riks

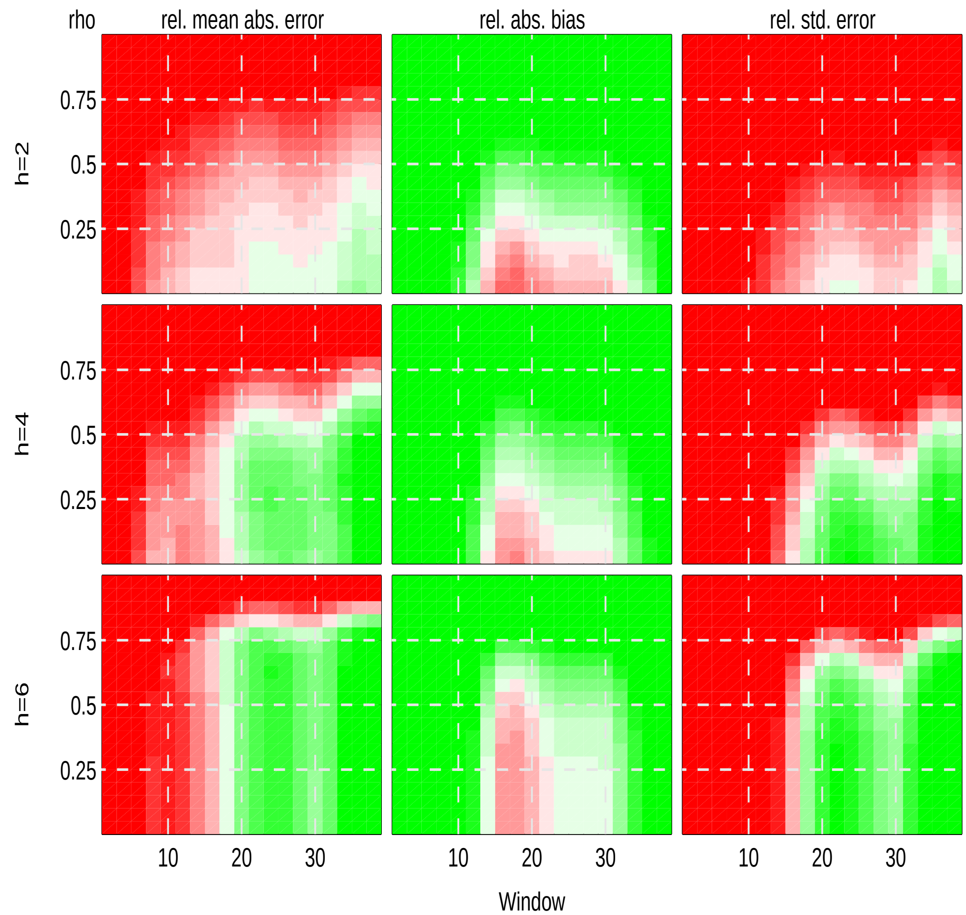

Fig. 2c: Performance of Output Forecast: SYSB(rho,window) relative to Riks